arXiv:quant-ph/0011036v1 9 Nov 2000

Quantum Information Theory

By

Michael Aaron Nielsen

B.Sc., University of Queensland, 1993
B.Sc. (First Class Honours), Mathematics, University of Queensland, 1994
M.Sc., Physics, University of Queensland, 1998

DISSERTATION

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy
Physics

The University of New Mexico
Albuquerque, New Mexico, USA

December 1998

©1998, Michael Aaron Nielsen

Dedicated to the memory of Michael Gerard Kennedy

24 January 1951 – 3 June 1998

Acknowledgments

It is a great pleasure to thank the many people who have contributed to this Dissertation. My deepest thanks goes to my friends and family, especially my parents, Howard and Wendy, for their support and encouragement. Warm thanks also to the many other people who have contributed to this Dissertation, especially Carl Caves, who has been a terrific mentor, colleague, and friend; to Gerard Milburn, who got me started in physics, in research, and in quantum information; and to Ben Schumacher, whose boundless enthusiasm and encouragement has provided so much inspiration for my research. Throughout my graduate career I have had the pleasure of many enjoyable and helpful discussions with Howard Barnum, Ike Chuang, Chris Fuchs, Raymond Laflamme, Manny Knill, Mark Tracy, and Wojtek Zurek. In particular, Howard Barnum, Carl Caves, Chris Fuchs, Manny Knill, and Ben Schumacher helped me learn much of what I know about quantum operations, entropy, and distance measures for quantum information. The material reviewed in chapters 3 through 5 I learnt in no small measure from these people. Many other friends and colleagues have contributed to this Dissertation. A partial list includes Chris Adami, Dorit Aharonov, James Anglin, David Beckman, Paul Benioff, Charles Bennett, Sam Braunstein, Nicholas Cerf, Ignacio Cirac, Richard Cleve, John Cortese, Vageli Coutsias, Wim van Dam, Ivan Deutsch, David DiVincenzo, Rob Duncan, Mark Ettinger, Betty Fry, Jack Glassman, Daniel Gottesman, Bob Griffiths, Gary Herling, Tom Hess, Alexander Holevo, Peter Høyer, Richard Hughes, Chris Jarzynski, Brad Johnson, Jeff Kimble, Alexei Kitaev, Andrew Landahl, Debbie Leung, Seth Lloyd, Hoi-Kwong Lo, Hideo Mabuchi, Eleanor Maes, Norine Meyer, Cesar Miquel, Juan Pablo Paz, Mitch Porter, John Preskill, Peter Shor, John Smolin, Alain Tapp, Barbara Terhal, Amnon Ta-Shma, Mike Westmoreland, Barby Woods, Bill Wootters, Christof Zalka, and Peter Zoller. My especial thanks also to those creative people without whom the physics of information would not exist as a field of research; especial influences on my thinking were the writings of Paul Benioff, Charles Bennett, Gilles Brassard, David Deutsch, Artur Ekert, Richard Feynman, Alexander Holevo, Richard Jozsa, Rolf Landauer, Göran Lindblad, Seth Lloyd, Asher Peres, Ben Schumacher, Peter Shor, John Wheeler, Stephen Wiesner, Bill Wootters, and Wojtek Zurek.

Michael Aaron Nielsen

The University of New Mexico
December 1998

Quantum Information Theory

By

Michael Aaron Nielsen

Abstract of Dissertation

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy
Physics

The University of New Mexico
Albuquerque, New Mexico, USA

December 1998

Quantum Information Theory

By

Michael Aaron Nielsen

B.Sc., University of Queensland, 1993
B.Sc. (First Class Honours), Mathematics, University of Queensland, 1994
M.Sc., Physics, University of Queensland, 1998

Abstract

Quantum information theory is the study of the achievable limits of information processing within quantum mechanics. Many different types of information can be accommodated within quantum mechanics, including *classical information*, *coherent quantum information*, and *entanglement*. Exploring the rich variety of capabilities allowed by these types of information is the subject of quantum information theory, and of this Dissertation. In particular, I demonstrate several novel limits to the information processing ability of quantum mechanics. Results of especial interest include: the demonstration of limitations to the class of measurements which may be performed in quantum mechanics; a capacity theorem giving achievable limits to the transmission of classical information through a two-way noiseless quantum channel; resource bounds on distributed quantum computation; a new proof of the quantum noiseless channel coding theorem; an information-theoretic characterization of the conditions under which quantum error-correction may be achieved; an analysis of the thermodynamic limits to quantum error-correction, and new bounds on channel capacity for noisy quantum channels.

Contents

List of figures **xi**

Nomenclature and notation **xiii**

I Fundamentals of quantum information **1**

1 The physics of information **3**
1.1 A collision of ideas . . . 3
1.2 What observables are realizable as quantum measurements? . . . 5
1.3 Overview of the field of quantum information . . . 8
1.4 Overview of the Dissertation . . . 9
1.5 Quantum information, science, and technology . . . 12

2 Quantum information: fundamentals **19**
2.1 Quantum bits . . . 19
2.2 Superdense coding . . . 21
2.3 Quantum teleportation . . . 23
2.4 Quantum computation . . . 26
2.5 What quantum measurements may be realized? . . . 28
2.6 Experimental quantum information processing . . . 30
2.6.1 Proposals for quantum information processing . . . 31
2.6.2 Experimental demonstration of quantum teleportation using NMR . . . 36

3 Quantum operations **41**
3.1 Quantum operations: fundamentals . . . 43
3.1.1 Quantum operations on a single qubit . . . 50
3.2 Freedom in the operator-sum representation . . . 51
3.3 Teleportation as a quantum operation . . . 54
3.4 Quantum process tomography . . . 57
3.4.1 One qubit example . . . 59
3.5 The POVM formalism . . . 62
3.6 Beyond quantum operations? . . . 63

4 Entropy and information **66**
4.1 Shannon entropy . . . 66
4.2 Basic properties of entropy . . . 67
4.2.1 The binary entropy . . . 67
4.2.2 The relative entropy . . . 68
4.2.3 Mutual information and conditional entropy . . . 69
4.2.4 The data processing inequality . . . 72
4.3 Von Neumann entropy . . . 73
4.3.1 Quantum relative entropy . . . 74
4.3.2 Basic properties of entropy . . . 75
4.3.3 Measurements and entropy . . . 76
4.3.4 The entropy of ensembles . . . 77
4.3.5 Subadditivity . . . 78
4.3.6 Concavity of the entropy . . . 79
4.4 Strong subadditivity . . . 80

5 Distance measures for quantum information **89**
5.1 Distance measures for classical information . . . 90
5.2 How close are two quantum states? . . . 92
5.2.1 Absolute distance . . . 92
5.2.2 Fidelity . . . 96
5.2.3 Distance measures derived from fidelity . . . 98
5.2.4 Relationships between distance measures . . . 99
5.3 Dynamic measures of information preservation . . . 100
5.3.1 Continuity relations . . . 104
5.3.2 Chaining quantum errors . . . 104
5.4 Alternative view of the dynamic measures . . . 105

II Bounds on quantum information transmission **108**

6 Quantum communication complexity **110**
6.1 The Holevo bound . . . 111
6.2 Capacity theorem for qubit communication . . . 116
6.3 Communication complexity of the inner product . . . 118
6.3.1 Converting exact protocols into clean form . . . 119
6.3.2 Reduction from the communication problem . . . 120
6.3.3 Lower bounds for bit protocols . . . 121
6.4 Coherent quantum communication complexity . . . 122
6.4.1 Coherent communication complexity of the quantum Fourier transform . . . 122
6.4.2 A general lower bound . . . 124
6.5 A unified model for communication complexity . . . 126
6.6 Conclusion . . . 128

7 Quantum data compression **130**
7.1 Schmidt numbers . . . 131
7.2 Typical subspaces . . . 133
7.3 Quantum data compression theorem . . . 136

7.4 Quantum data compression with a classical side channel . . . 139
7.5 Universal data compression with a classical side channel . . . 141
7.5.1 A dense subset of density operators with distinct entropies . . . 143
7.6 Conclusion . . . 146

8 Entanglement **147**
8.1 Pure state entanglement . . . 148
8.2 Mixed state entanglement . . . 149
8.3 Entanglement: Examples . . . 152
8.4 Conclusion . . . 161

9 Error correction and Maxwell's demon **163**
9.1 Entropy exchange . . . 164
9.2 Quantum Fano inequality . . . 166
9.3 The quantum data processing inequality . . . 167
9.4 Quantum error correction . . . 169
9.4.1 Shor's code . . . 173
9.5 Information-theoretic conditions for error correction . . . 175
9.6 Information-theoretic inequalities for quantum processes . . . 178
9.7 Quantum error correction and Maxwell's demon . . . 181
9.7.1 Error-correction by a "Maxwell demon" . . . 181
9.7.2 Discussion . . . 186
9.8 Conclusion . . . 186

10 The quantum channel capacity **188**
10.1 Noisy channel coding . . . 189
10.2 Classical noisy channels in a quantum setting . . . 190
10.3 Coherent information . . . 192
10.3.1 Properties of coherent information . . . 192
10.3.2 The entropy-fidelity lemma . . . 194
10.3.3 Quantum characteristics of the coherent information I . . . 194
10.3.4 Quantum characteristics of the coherent information II . . . 196
10.4 Noisy channel coding revisited . . . 197
10.4.1 Mathematical formulation of noisy channel coding . . . 198
10.5 Upper bounds on the channel capacity . . . 200
10.5.1 Unitary encodings . . . 200
10.5.2 General encodings . . . 203
10.5.3 Other encoding protocols . . . 205
10.6 Discussion . . . 206
10.7 Channels with a classical observer . . . 207
10.7.1 Upper bounds on channel capacity . . . 209
10.7.2 Relationship to unobserved channel . . . 212
10.7.3 Discussion . . . 213
10.8 Conclusion . . . 213

III Conclusion **217**

11 Conclusion **219**
11.1 Summary of the Dissertation 219
11.2 Open problems in quantum information 221
11.2.1 A unifying picture for quantum information 222
11.2.2 Classical physics and the decoherence program 224
11.2.3 Quantum information and statistical physics 225
11.3 Concluding thought 226

A Purifications and the Schmidt decomposition **227**

Bibliography **231**

Index **244**

List of figures

1.1 Adaptation of the meaning circuits proposed by Wheeler [186] and Landauer [107]. 15

2.1 Circuit for quantum teleportation. The measurement is in the computational basis, leaving the measurement result stored in the data and ancilla qubits. R_y and R_{-y} denote rotations of 90 degrees about the y and $-y$ axes on the Bloch sphere. . . . 24

2.2 Schematic representation of the labeled TCE molecule. The two carbon atoms, C1 and C2, are ^{13}C isotopes, which a net nuclear spin of 1/2. 32

2.3 Schematic of quantum teleportation. See text for a full description. 37

2.4 Dynamic fidelity is plotted as a function of decoherence time. The bottom curve is a control run where the information remains in C2. The curve shows a decay time of approximately 0.5s. The top curve represents the fidelity of the quantum teleportation process. The decay time is approximately 2.6s. The information is preserved for a longer time, corresponding approximately to the combined effects of decoherence and relaxation for the Hydrogen, confirming the prediction of teleportation. 39

3.1 Environmental model for a quantum operation. 48

4.1 Binary entropy function $H(p)$. 67

4.2 Relationships between different entropies. 71

5.1 The RQ picture of a quantum channel. The initial state of RQ is a pure state. . . 101

6.1 Plot of the Holevo bound when the states $|0\rangle$ and $\cos\theta|0\rangle + \sin\theta|1\rangle$ are prepared with equal probability. Notice that the Holevo bound reaches a maximum when the angle between the two states is $\theta = \pi/2$, corresponding to orthogonal states. At this point only it is possible for Bob to determine with certainty which state Alice prepared. 114

6.2 Capacity region to send n bits from Alice to Bob. n_{AB} is the number of qubits Alice sends to Bob, and n_{BA} is the number of qubits Bob sends to Alice. The dashed line indicates the bottom of the classical capacity region. 116

7.1 Quantum data compression. The compression operation $\mathcal{C}$ compresses a quantum source ρ stored in $n \log d$ qubits into $nS(\rho)$ qubits. The source is accurately recovered via the decompression operation $\mathcal{D}$. 130

7.2 Quantum data compression with a classical side channel. Results of measurements made during compression are made available to assist during decompression. . . . 140

7.3 Construction of a data compression scheme which works on all elements in π. The vertical columns indicate a compression scheme that is optimal with respect to the set of source density operators at the head of the column. For example, the second column contains a compression scheme which is optimal with respect to both ρ_1 and ρ_2. From this table the underlined elements are used to construct a new compression scheme which is optimal with respect to all elements of π. 144

8.1 Entanglement (E) plotted as a function of temperature (T) for a coupling strength $b = 2$ in the regime where the coupling dominates. 155
8.2 Entanglement (E) plotted as a function of temperature (T) for a weak coupling strength, $b = 0.9$. 156
8.3 Entanglement (E) plotted as a function of the coupling strength (b) at zero temperature. 157
8.4 Entanglement (E) plotted as a function of the coupling strength (b) at a finite temperature, $T = 1$. 158
8.5 Entanglement (E) plotted as a function of temperature (T) and the coupling strength (b). This figure shows the behaviour for values of the coupling strength near the crossover at $b = 1$. 159
8.6 Entanglement (E) plotted as a function of temperature (T) and the coupling strength (b). This figure shows the behaviour for large values of the coupling strength, $b \gg 1$. 160

9.1 Binary symmetric channel. 170
9.2 Encoding circuits for the (a) bit flip and (b) phase flip codes. 173
9.3 Encoding circuit for the Shor nine qubit code. 174
9.4 Error correction cycle. 183

10.1 The noisy quantum channel, together with encodings and decodings. 189
10.2 Dual classical channels operating on inputs X_1 and X_2 produce outputs Y_1 and Y_2. 196
10.3 The noisy quantum channel, together with encodings and decodings. 198
10.4 Noisy quantum channel with an extra stage, a restricted *pre-encoding*, $\mathcal{P}$. 207
10.5 Noisy quantum channel with a classical observer. 208

Nomenclature and notation

There are several items of nomenclature and notation which have two or more meanings in common use in the field of quantum information theory. To prevent confusion from arising, this section collects many of the more frequently used of these items, together with the conventions that will be adhered to in this Dissertation.

As befits good information theorists, logarithms are *always* taken to base two, unless otherwise noted.

A *positive* operator A is one for which $\langle\psi|A|\psi\rangle \geq 0$ for all $|\psi\rangle$. A *positive definite* operator A is one for which $\langle\psi|A|\psi\rangle > 0$ for all $|\psi\rangle \neq 0$.

The *relative entropy* of a positive operator A with respect to a positive operator B is defined by

$$S(A||B) \equiv \mathrm{tr}(A \log A) - \mathrm{tr}(A \log B). \tag{1}$$

Conventionally, most researchers use $|\psi\rangle$ to represent a pure state of a quantum system, and ρ to represent a mixed state. We will use this notation on occasion, but we will also make use of a different notation. Suppose we are dealing with a composite quantum system with component parts labeled R and Q. Then we will use R, Q, and RQ to denote the quantum states associated with those systems, in addition to their use as labels for the systems. When one or more of these systems is known to be in a pure state we will use the notation $|R\rangle$, $|Q\rangle$, and $|RQ\rangle$, as appropriate.

A *purification* of a mixed state, Q, of some quantum system Q, is a pure state $|RQ\rangle$ of some larger system RQ, such that when the system R is traced out, the state Q is recovered,

$$Q = \mathrm{tr}_R(|RQ\rangle\langle RQ|). \tag{2}$$

The *support* of an operator is defined to be the vector space orthogonal to its kernel. For a Hermitian operator, this means the vector space spanned by eigenvectors of the operator with non-zero eigenvalues.

The term *probability distribution* is used to refer to a finite set of real numbers, p_x, such that $p_x \geq 0$ and $\sum_x p_x = 1$.

All Hilbert spaces are assumed to be finite dimensional. In many instances this restriction is unnecessary, or can be removed with some additional technical work, but making the restriction globally makes the presentation more easily comprehensible, and doesn't detract much from many of the intended applications of the results. Furthermore, in some instances, extension of a result to general Hilbert spaces is beyond my technical expertise!

Part I

Fundamentals of quantum information

Chapter 1

The physics of information

> Information is physical.
>
> – Rolf Landauer [108]

1.1 A collision of ideas

What is discovered when the laws of physics are used as the foundation for investigations of information processing and computation? This Dissertation is an attempt to provide a partial answer to this question. To understand why the attempt should be fruitful, it is useful to remind ourselves of what it is that a physicist or computer scientist does.

What is physics? Physics is a messy human endeavour, so any answer to this question is somewhat inaccurate. Nevertheless, an examination of the history and current state of physics reveals at least two overarching themes within physics. The first theme is that physics studies *universal* properties of nature. We expect that black holes in the cores of galaxies a billion light years away obey the same laws of general relativity that govern the motion of planets in our own solar system, or the motion of a ball through the air. Likewise, we expect that the structure of matter based upon quarks and leptons is the same throughout the universe.

A second overarching theme of physics is the *reduction of phenomena*. This theme has two aspects. One aspect is the ongoing search for simplified, unified frameworks in which it is possible to understand more complicated phenomena. For example, there is the current search for a unified description of the particles and fields of nature [191], or attempts to understand the principles underlying pattern formation in physics [109]. The second aspect of this theme is the discovery and explanation of phenomena in terms of simple frameworks. For example, there is the remarkable Bardeen-Cooper-Schrieffer theory of superconductivity [6], based upon the principles of quantum mechanics, or the current search for gravitational waves [175], potentially one of the most useful consequences of the general theory of relativity.

Note that both these themes are somewhat gray. There are differing degrees of universality, and physics does not concern itself with the reduction of all phenomena to fundamentals. It leaves many phenomena – the human body, climate patterns, computer design – to other disciplines. Here too, universality plays a role, with physics being primarily interested in relatively simple phenomena, such as superfluidity, which do not have an especially detailed historical dependence such as may be found, for example, in the functioning of a cell, and can therefore be relatively easily reproduced by a variety of means, in many locations.

The themes of universality and reduction both have strong parallels within computer science[1]. Traditionally, computer science is based upon a small number of *universal* models that are each supposed to capture the essence of some aspect of information processing. For example, the majority of work done on algorithm design has been framed within the well known Turing machine model [177] of computation, or one of its equivalents. Shannon's model [162] of a communications channel is the foundation for modern work in information theory.

Computer science is also concerned with the reduction of phenomena, but in a different way than is often the case in physics. Reduction in physics often concerns the explanation of phenomena discovered without specific intent, such as superconductivity. In computer science, it is more typical to set a specific information processing goal – "I would like my computer to sort this list of names for me in such and such an amount of time" – and then to attempt to meet that goal within an existing model of information processing.

What is the origin of the fundamental models used as the basis for further progress in computer science? Examination of the original papers shows that the founders used systems existing in the real world as inspiration and justification for the models of computation they proposed. For example, Turing analyzed the set of operations which a mathematician could perform with pen and paper, in order to help justify the claim that his model of computation was truly universal.

It is a key insight of the last thirty years that these pseudophysical justifications for the fundamental models of computation may be carried much further. For example, a theory of computation which has its foundations in quantum mechanics has been formulated [63]. *Information is physical*, as Landauer reminds us [108]. That is, any real information processing system relies for its implementation upon systems whose behaviour is completely described by the laws of physics.

Remarkable progress has been achieved by acting on this insight, re-examining and reformulating the fundamental models of information based upon physical principles. The hope, which has been fulfilled, is that such a reformulation will reveal information processing capabilities that go beyond what was thought to be possible in the old models of computation.

The field of science which studies these fundamental connections between physics and information processing has come to be known as the *physics of information*. The connection between physics and information processing is a two way street, with potential benefits for both computer science and physics.

Computer science benefits from physics by the introduction of new models of information processing. *Any physical theory may be regarded as the basis for a theory of information processing.* We may, for example, enquire about the computational power of Einstein's general theory of relativity, or about the computational power of a quantum field theory. The hope is that these new models of information processing may give rise to capabilities not present in existing models of information processing. In this Dissertation we will primarily be concerned with the information processing power of quantum mechanics. The other possible implication for computer science is more ominous: there may be unphysical elements in existing theories of information processing which need to be rooted out if those theories are to accurately reflect reality.

Physics benefits in at least four ways from computer science. First, computer science may act as a stimulus for the development of new techniques which assist in the pursuit of the fundamental goals of physics. For example, inspired by coding theory, error correction methods to protect against noise in quantum mechanics have been developed. One of the chief obstacles to precision measurement is, of course, the presence of noise, so error correcting codes to reduce the effects of that noise are welcome. They are doubly useful, however, as a diagnostic tool, since error correcting codes can be used to determine what types of noise occur in a system.

[1] I am using "computer science" as a pseudonym for all those fields of science concerned with information processing, including, for example, computer science, information theory, signal processing, and many others.

The second way physics benefits from computer science is via simulation. Computational physics has allowed us to investigate physical theories in regimes that were not previously accessible. Such investigations can lead to interesting new questions about those theories, and yield important insights into the predictions made by our physical theories.

The third way physics benefits from computer science is that computers enable us to perform experiments that would once have been impossible or, at the least, much more difficult and expensive. Computer-based methods for obtaining, analysing, and presenting data have opened up new experimental realms. For example, computers enormously simplify the analysis of data taken in particle accelerators, in which only a miniscule fraction of the events detected in a given experimental run may be of direct interest. Automated sifting of the data and identification of the relevant events is performed in an instant using powerful computers, rather than the time of years or more that it would take a human being to achieve the same results.

The fourth way physics benefits from computer science is more difficult to describe or justify. My experience has been that computer science is a great inspiration for fundamental questions about physics, and can sometimes suggest useful approaches to take in the solution of physics problems. This will be apparent several times during the main body of this Dissertation. I can not yet say precisely why this should be the case, although as we have seen, both physics and computer science involve the development of tools to reduce phenomena involving complex interacting systems to certain fundamental models, as well as continual questioning and refinement of those models. Perhaps it is not so surprising that each field should have much to teach the other.

1.2 What observables are realizable as quantum measurements?

This Dissertation is concerned principally with a special subfield of the physics of information, *quantum information*, in which the fundamental models for information processing are based upon the laws of quantum mechanics. The earlier formulation of the question investigated by this Dissertation may thus be refined: *What is discovered when the laws of quantum mechanics are used as the foundation for investigations of information processing and computation?*

To better understand the subject of quantum information, it is useful to have a concrete example in hand. This section presents a simple example which illustrates many of the basic themes of quantum information. The example is also interesting in its own right, as it takes us straight to the edge of what is known, posing a fundamental question about quantum mechanics, inspired by the methods of computer science.

The example concerns the question of *what properties of a quantum mechanical system may be measured?* In the 1920s, Heisenberg and other researchers formulated the notion of a *quantum mechanical observable*. Observables were introduced into quantum mechanics as a means of describing what properties of a quantum system may be measured. For example, a particle's position is regarded as an observable in quantum mechanics.

Mathematically, the concept of an observable is usually formulated as follows. An observable is any Hermitian operator acting on the state space of a physical system, where by "state space" we shall mean the usual Hilbert space associated with a physical system. Recall from elementary quantum mechanics that the measurement postulate of quantum mechanics as usually formulated has the following consequences: To each measurable quantity of a quantum mechanical system there is associated a mathematical object, an *observable*, which is a Hermitian operator acting on the state space of the quantum system. The possible outcomes of the measurement are given by the spectrum of the observable. If the state of the quantum system immediately before the system is

observed is an eigenstate of the observable then, with certainty, the outcome of the measurement is the corresponding eigenvalue, m.

One of the most remarkable discoveries of quantum mechanics is that the theory implies limits to the class of measurements which may be performed on a physical system. The most famous example of this is the *Heisenberg uncertainty principle*, which establishes fundamental limits upon our ability to perform simultaneous measurements of position and momentum. Given the shock caused by Heisenberg's result that there are limits, in principle, to our ability to make observations on a physical system, it is natural to ask for a precise characterization of what properties of a system may be measured. For example, Dirac's influential text ([61], page 37) makes the following assertion on the subject:

The question now presents itself – Can every observable be measured? The answer theoretically is yes. In practice it may be very awkward, or perhaps even beyond the ingenuity of the experimenter, to devise an apparatus which could measure some particular observable, but the theory always allows one to imagine that the measurement can be made.

That is, Dirac is asserting that given any observable for a reasonable quantum system, it is possible in principle to build a measuring device that makes the measurement corresponding to that observable. Dirac leaves his discussion of the subject at that, making no attempt to further justify his claims. Later, Wigner [190] investigated the problem, and discovered that conservation laws do, in fact, impose interesting physical constraints upon what properties of a system may be measured. This work was subsequently extended by Araki and Yanase [5], resulting in what Peres [140] terms the Wigner-Araki-Yanase or WAY theorem. To my knowledge, there has been remarkably little other work done on the fundamental question of what observables may be measured in quantum mechanics.

Not long after Heisenberg, Dirac and others were laying the foundations for the new quantum mechanics, a revolution of similar magnitude was underway in computer science. The remarkable English mathematician Alan Turing laid out the foundations for modern computer science in a paper written in 1936 [177][2].

Turing's work was motivated, in part, by a challenge set down by the great mathematician David Hilbert at the International Congress of Mathematicians held in Bologna in 1928. Hilbert's problem, the *Entscheidungsproblem*, was to find an algorithm by which all mathematical questions could be decided. Remarkably, Turing was able to show that there is *no such procedure*. Turing demonstrated this by giving an explicit example of an interesting mathematical question whose answer could not be decided by algorithmic means. In order to do this, Turing had to formalize our intuitive notion of what it means to perform some task by algorithmic means.

To do this, Turing invented what is now known as the *universal Turing machine*. Essentially, a universal Turing machine behaves like an idealized modern computer, with an infinite memory. Turing's computer was capable of being *programmed*, in much the same sense as a modern computer may be programmed. Turing's programs computed mathematical functions: the machine would take a number as input, and return a number as output, with the function computed by the machine in this way being determined by the program being run on the machine. In addition, it was possible that programs would fail to halt, continuing to execute forever, never giving a definite output.

The most important assertion in Turing's paper has come to be known as the *Church-Turing thesis*. Roughly speaking this thesis states that *any function which may be computed by what we intuitively regard as an algorithm may be computed by a program running on a universal Turing machine, and vice versa*. The reason this thesis is so important is because it asserts the equivalence

[2]It is worth noting that many other researchers arrived at similar results around the same time, notably Church and Post. However it is my opinion that it is Turing's grand vision that has ultimately proved to be the deepest and most influential.

of an intuitive concept – that of an algorithm – with the rigorously defined mathematical concept of a program running on a universal Turing machine. The validity of the Church-Turing thesis has been repeatedly tested and verified inductively since Turing's original paper, and it is this continuing success that ensures that Turing's model of computation, and others equivalent to it, remain the foundation of theoretical work in computer science.

One observation made by Turing was that the programs for his universal machine could be numbered, $0, 1, 2, \ldots$. This led him to pose the *halting problem*: does program number x halt on input of the value x, or does it continue forever? Turing showed that this apparently innocuous question has no solution by algorithmic means. In fact, it is now known that in some sense "most" questions admit no algorithmic solution. The way Turing demonstrated the unsolvability of the halting problem was to note that it is equivalent to being able to compute the *halting function*,

$$h(x) \equiv \begin{cases} 1 & \text{if program } x \text{ halts on input } x \\ 0 & \text{if program } x \text{ does not halt on input } x, \end{cases} \tag{1.1}$$

by algorithmic means.

In Chapter 2 we review the proof that there is no algorithm which can compute the halting function, establishing Turing's great result. For now, we will assume that this remarkable result is correct.

Turing's result paves the way for an interesting quantum mechanical construction. Suppose we consider a quantum mechanical system whose state space is spanned by orthonormal states $|0\rangle, |1\rangle, \ldots$, such as the quantum mechanical simple harmonic oscillator. We use the halting function to define a Hermitian operator, $\hat{h}$, by the formula:

$$\hat{h} \equiv \sum_{x=0}^{\infty} h(x)|x\rangle\langle x|. \tag{1.2}$$

This operator is clearly Hermitian, and thus represents a quantum mechanical observable, which we call the *halting observable*. Notice that it has two eigenvalues, 0 and 1. The eigenspace corresponding to the eigenvalue 1 is spanned by those states $|x\rangle$ for which $h(x) = 1$, while the eigenspace corresponding to the eigenvalue 0 is spanned by those states $|x\rangle$ for which $h(x) = 0$.

Is the halting observable a measurable property of the quantum mechanical system? More precisely, is it possible to construct a measuring device which performs a measurement of the halting observable? There are two possibilities:

1. It is possible, in principle, to construct a measuring device which can measure the halting observable. In this case, we can give a physical algorithm for solving the halting problem: to evaluate $h(x)$, build the device to measure the halting observable, prepare the quantum system in the state $|x\rangle$, and perform a measurement of the halting observable. By the quantum measurement postulate, the result of the measurement is, with certainty, the correct value of $h(x)$.

2. It is not possible, in principle to construct a measuring device which can measure the halting observable.

If the first possibility is correct, we are forced to conclude that Turing's model of computation is insufficient to describe all possible algorithms, and thus the Church-Turing thesis needs to be re-evaluated. If, on the other hand, the second possibility is correct, then we are left to ponder the problem of determining the fundamental limits to measurement in quantum mechanics.

A resolution of this dichotomy, which I posed in [132][3], is not presently known. It is instructive to note several features of the problem posed. First, it is a problem concerning the ultimate limits to our ability to perform a particular "information processing" task, in this case, the performance of a quantum measurement. We are interested in finding limitations on what is possible, and also in constructive techniques for certain physical tasks. In the present example of measurement theory, I think it is fair to say that we do not yet fully understand either the limits or the possibilities available in the measurement process. Second, it is interesting to note the fruitful interplay between physics and computation taking place here: a fundamental question from computer science has been translated into physical terms, and gives rise to an interesting fundamental question about physics. Both these features are repeated many times through the course of this Dissertation, and throughout quantum information in general.

With this concrete example in hand, we now return to understand in more detail what the subject of quantum information is about. In Chapter 2 we return to study the problems posed by the halting observable and similar constructions in greater depth.

1.3 Overview of the field of quantum information

Quantum information may be defined as *the study of the achievable limits to information processing possible within quantum mechanics*. Thus, the field of quantum information has two tasks.

First, it aims to determine limits on the class of information processing tasks which are possible in quantum mechanics. For example, one might be interested in limitations on the class of measurements that may be performed on a quantum system – if it is impossible to measure the halting observable, then that would be an interesting fact to know, and explore in greater detail. Another example which will be examined in this Dissertation is the question of determining bounds to how much information may be stored using given quantum resources.

The second task of quantum information theory is to provide constructive means for achieving information processing tasks. For example, it would be extremely useful to have a means for implementing any desired measurement in quantum mechanics. Another example, where this goal of constructive success has to some extent been achieved, is in the development of unbreakable schemes for doing *cryptography*, based upon the principles of quantum mechanics [188, 16, 88]. This is an especially interesting example, as Shannon used the tools of classical information theory to "prove" that the task accomplished by quantum cryptography was not possible [161]. Of course, the flaw in Shannon's proof is that he assumed a model of communication that did not include the possibilities afforded by quantum mechanics.

Ideally, these two tasks would dovetail perfectly; for each limit to information processing that we prove, we would find a constructive procedure for achieving that limit. Alas, that ideal is often not achieved, although it remains a central goal of all investigations into quantum information processing.

We have been rather vague about what is meant by the term "quantum information". What sorts of entities qualify as quantum information? The answer to this question will evolve as we proceed through the Dissertation, however it is useful to look ahead at what is in store, by taking a historical tour to enlighten us as to the tasks which may be performed using quantum information.

Quantum information really began to get going during the 1960s and 1970s. For example, several researchers began to ask and answer questions about what communications tasks could be accomplished using quantum states as intermediary resources. The inputs and outputs to the processes considered were usually classical information, with the novelty coming from the use of

[3] Benioff (private communication) has independently constructed related examples, with similar ends in mind.

quantum resources during the process to aid in the accomplishment of the task. The questions being asked about these processes were framed in terms of classical information processing. Much of this early work is reviewed in the inspiring books of Holevo [84] and Helstrom [81].

A little later came the invention of quantum cryptography [188, 17] and the quantum computer [13, 66, 58]. In these applications the role of quantum mechanics is rather more subtle. Both possibilities, most decisively quantum cryptography, enable the performance of information processing tasks which are considered "impossible" in classical information theory. These new information processing capabilities acted as a great motivator for the idea that essentially new types of "information" were being used to perform these tasks – quantum information. Moreover, in both applications it is necessary to take into account the effect of noise on quantum states, and if possible minimize the effect of that noise. That problem is strongly reminiscent of the problem of protecting against noise which arises in classical information theory, yet without any classical "information" apparently involved in the process.

More recently, the pioneering work of Schumacher [154] on quantum data compression, by Shor [164] and Steane [172] on quantum error correcting codes, and by Wootters and coworkers [20, 22, 82, 195] on measures of entanglement has made this idea of essentially new types of information much more precise. Schumacher quantified the physical resources necessary to store the quantum states being emitted by a "quantum source". Shor and Steane showed how to protect quantum states and entanglement against the effects of noise. Finally, Wootters and coworkers have emphasized the use of quantum entanglement as a resource that may be useful in the solution of many information processing problems, and have characterized entanglement by its efficiency as an aid in those problems.

Perhaps, then, we may distill the following heuristic definition of information from this historical tour: *(Quantum) Information is any physical resource which may be of assistance in the performance of an interesting (quantum) information processing task.* Of course, this simply moves the definitional difficulty elsewhere, but speaking for myself, I believe that I have a better intuitive feel for what constitutes an information processing task than for the more ethereal question of what information is.

Reflecting on this historical tour we see that quantum information comes in many different types. Some of the types of information of interest include classical information, entanglement, and actual quantum states. This is in contrast to the classical theory of information processing which is largely focused on information types derived from a single structure: the bit[4]. The greater variety of information structures available in quantum information necessitate a broader range of tools for understanding the different information types, and open up a richer range of information processing possibilities for exploration.

We conclude from the history that quantum information is an evolving concept, and it seems likely that we are yet a long way from grasping all the subtleties of the different kinds of quantum information. Indeed, in the future we may discover new quantum resources whose importance is not yet glimpsed, but which will one day be seen as a crucial part of quantum information theory.

1.4 Overview of the Dissertation

The primary purpose of the Dissertation is *to develop theoretical bounds on our ability to perform information processing tasks in quantum mechanics.*

[4]There is a well developed theory of analogue computation, which, however, appears to be equivalent to the theory based on bits, when physically realistic assumptions about the presence of noise are made.

Two aspects of this purpose deserve special comment. First, it is to be emphasized that the purpose is to find bounds on our ability to process information. We will not always be able to determine whether the bounds we discover are achievable. Nevertheless, it is still of considerable interest to understand limits to what is in principle possible. Second, the focus of the Dissertation is theoretical, although Chapter 2 does contain an overview of the experimental state of the field, and the results of a simple experiment in quantum information. However, a full exposition of the experimental state of the field is beyond the scope of this Dissertation.

The Dissertation is structured into three parts.

The first part of the Dissertation, "Fundamentals of quantum information", provides an introductory overview of quantum information, and develops tools for the study of quantum information. Part I consists of Chapters 1 through 5. The primary purpose of Part I is to provide a pedagogical introduction and reference for concepts in the field. While Part I contains a substantial amount of original research material, the presentation of that material is ancillary to the main goal, which is to provide a solid basis for the understanding of the quantum information-theoretic problems investigated in Part II of the Dissertation, which is primarily oriented towards original research results.

The following is a brief summary of the contents of Part I.

Chapter 2 provides an introduction to many of the most basic notions used in quantum information, such as quantum states, dynamics, quantum gates, quantum measurements and the notion of a quantum computer. These notions are illustrated using a number of simple examples, most notably quantum teleportation and superdense coding. We revisit in greater detail the question of what measurements may be performed in a quantum system. The Chapter concludes with a summary of some of the challenges facing experimental quantum information. Notable original features of the Chapter include a discussion of realizable measurements in quantum mechanics, and the description of an experimental implementation of quantum teleportation using nuclear magnetic resonance.

Chapter 3 is a review of the *quantum operations* formalism, used to describe state changes in quantum systems. This formalism includes as special cases the unitary evolution generated by the Schrödinger equation, quantum measurements, and noise processes such as phase decoherence and dissipation. Notable original features of the Chapter include a discussion of *quantum process tomography*, a procedure by which the dynamics of a quantum system may be experimentally determined, and a formulation of quantum teleportation within the quantum operations formalism.

Chapter 4 reviews the concepts of *entropy* and *information* that underpin much of quantum information. Entropic measures often arise naturally in the study of resource problems in quantum information, which are usually of the form *how much of physical resource X do I need to accomplish task Y?* Much of Part II of the Dissertation is concerned with such resource problems, so it is crucial that we obtain a solid understanding of the basic facts about entropy. A notable feature of the Chapter is the inclusion of several inequalities relating von Neumann entropies which I believe to be new.

Chapter 5 reviews *distance measures for quantum information*. A distance measure provides a means for determining the similarity of two items of quantum information. For example, we may be interested in the question of what it means for two quantum states to be "close" to one another. Many different measures of distance may be proposed, motivated by different physical questions one may ask about quantum information. This Chapter reviews the motivation for many of these definitions, and attempts to relate some of the definitions that have been proposed. This Chapter contains the most original material of any Chapter in Part I, including many new properties of the various measures of distance investigated, and some new relationships between the distance measures that have been proposed.

This concludes the summary of the contents of Part I.

Part II of the Dissertation, "Bounds on quantum information transmission", poses a number of questions about information transmission, and provides bounds on the answers to those questions. Part II consists of Chapters 6 through 10 of the Dissertation. In Part II, the tools developed in Part I are employed in the investigation of several substantive questions in quantum information theory. Part II is largely devoted to the presentation of original research results.

The following is a brief summary of the contents of Part II.

Chapter 6 studies quantum communication complexity. Quantum communication complexity is concerned with the communication cost incurred during the performance of some distributed computation, if quantum resources are employed for the communication. Recently, several remarkable results have been proved showing that in some cases the use of quantum resources may provide a substantial saving over the communication cost required to solve a problem in distributed computation with classical resources. The Chapter begins with an explanation of Holevo's theorem, which is a fundamental bound on the ability to perform quantum communication. This bound is then applied to give a new capacity theorem which precisely quantifies the resources required to send classical information over a two way quantum noiseless channel. This capacity theorem is applied to demonstrate a significant new negative result in quantum information: that there exist problems of distributed computation for which the use of quantum resources can provide no improvement over the situation in which only classical resources are used. Next, we turn our attention to the following problem: what communication resources are required to compute a *quantum function* – a unitary evolution – if that function is distributed over two or more parties? To my knowledge all previous work on quantum communication complexity has focused on distributed computation of *classical functions*. In addition to posing this problem for the first time, this Chapter contains the first non-trivial lower-bound on such a problem, the communication complexity for computation of the quantum Fourier transform by two parties, as well as a general lower bound for the communication complexity of an arbitrary unitary operator.

Chapter 7 studies the problem of quantum data compression. It is well known that it is often possible to compress classical information so that it uses up fewer physical resources. For example, there are many widely used programs which can be used to compress computer files so that they take up less disk space. It turns out that it is possible to compress quantum states along somewhat similar lines, so that they may be stored using fewer physical resources. This Chapter provides a new proof of the fundamental theorem of quantum data compression, substantially simplifying earlier proofs. Furthermore, the Chapter reports results on *universal data compression*, which allows the compression of a quantum source whose characteristics are not completely known.

Chapter 8 studies the fundamental problem of providing quantitative measures of the entanglement between two quantum systems. More than any other resource, it appears to be quantum entanglement which enables the most striking departures of quantum information processing from classical information processing, and it is to be hoped that developing quantitative measures of entanglement will enable us to better understand the nature of this resource. Several measures of entanglement are reviewed, and many new bounds on these measures and relationships between the measures are proved. I discuss the insights into quantum information which are given by these bounds, emphasizing connections with other problems studied in the Dissertation. This Chapter does not contain any results which are especially striking in their own right; rather it proves several new results and examples which provide insight into the results in other Chapters of Part II.

Chapter 9 describes the methods that have been developed for the performance of quantum error correction. New information-theoretic conditions for quantum error-correction are developed, together with other information-theoretic constraints upon the error-correction process. The Chapter concludes with an original analysis of the thermodynamic cost of quantum error correction.

Chapter 10 studies the *capacity* of a noisy quantum channel. The capacity is a measure of how much quantum information can be transferred over a noisy quantum communications channel with high reliability. Unfortunately, at present the quantum channel capacity is still rather poorly understood. This Chapter presents several new bounds on the quantum channel capacity, and emphasizes the differences between classical and quantum information which make the quantum channel capacity especially interesting. The Chapter concludes with the presentation of a new problem in quantum information theory, that of determining the quantum channel capacity of a noisy quantum channel in which partial classical access to the channel environment is allowed. New expressions upper bounding the capacity in this instance are proved.

This concludes the summary of the contents of Part II.

Part III of the Dissertation, "Conclusion", consists of a single Chapter, Chapter 11, which summarizes the results of the Dissertation, and sketches out some directions for future work. Chapter 11 begins with a brief summary of the results of the Dissertation, highlighting specific questions raised in the Dissertation which deserve further investigation. The Chapter and main body of the Dissertation concludes by taking a broader look at the future directions available to quantum information theory, outlining a number of possible research programs that might be pursued.

Some miscellaneous remarks on the style and structure of the Dissertation:

The front matter of the Dissertation contains a detailed table of contents, which I encourage you to read, as well as a list of figures with their associated captions. There is also a guide to nomenclature and notation, which contains notes to assist the reader in translation between the often incompatible conventions used by different authors in the field of quantum information.

Each Chapter in Parts I and II of the Dissertation begins with an overview of the problems to be addressed in the Chapter, and concludes with a boxed summary of the main results of the Chapter. Collaborations with other researchers are indicated where appropriate, usually at the beginning of a Chapter or section. In addition, I have tried whenever possible to give credit for prior work in the field, with citations pointing to the extensive bibliography which may be found at the end of the Dissertation. My apologies to any researcher whose work I have inadvertently omitted. Ike Chuang supplied figures 4.2, 9.2 and 9.3.

The end matter of the Dissertation contains a single Appendix, a Bibliography, and an Index.

The Appendix contains material which I felt was outside the main thrust of the Dissertation, but nevertheless is sufficiently interesting and useful to warrant inclusion. It discusses the Schmidt decomposition, a structural theorem useful for the study of composite quantum systems. A new generalization of the Schmidt decomposition is proved in the Appendix, and related concepts such as purifications of mixed states are discussed.

The Bibliography contains a listing of all reference materials cited in the text of the Dissertation, ordered alphabetically by the family name of the first author.

The Index references the most important occurrences of technical terms and results appearing in this Dissertation. Only subjects are indexed, not names.

Finally, I note that the Dissertation has been written in the first person. When "I" appears, it indicates my opinion, or something for which I claim responsibility. "We" indicates occasions where I hope you, the reader, and I, the author, can fully agree.

1.5 Quantum information, science, and technology

What is the broader relationship of quantum information with science and technology? This question is well beyond my ability to answer in full, however based upon what we now know it is interesting to essay some possible answers.

Let us start with science. Predicting the future impact of quantum information on science is obviously impossible in detail (although see Chapter 11 for an attempt in this direction). Instead, we will attempt to relate the existing goals and achievements of quantum information to other areas of science, and science as a whole.

An area in which quantum information theory has already had a substantial impact is on physicists' understanding of quantum mechanics. Quantum mechanics is legendary for the counter-intuitive nature of its predictions. One way to lift the veil of mystery surrounding quantum mechanics is to develop a toolkit containing simple tools on which we can rely to help us navigate quantum mechanics. The development of such a toolkit is one of the primary aims of quantum information, and is the central goal of Part I of this Dissertation.

One consequence of this tool building is the development of many equivalent ways of formulating fundamental physical principles. For example, Westmoreland and Schumacher [185] have recently argued that the physical prohibition against superluminal communication can be deduced from elementary quantum mechanics, via the no-cloning theorem [196, 60]. Feynman [65] has argued that such development of new ways of looking at physical principles has great value for fundamental research. As we do not yet have a complete fundamental physical theory of the world [191, 193], new perspectives on old theories such as quantum mechanics may be extremely useful in the search for a more complete theory of the world.

A second area in which I expect the physics of information to eventually have a great impact is in the study of statistical physics and collective phenomena. Collective phenomena involve large numbers of systems interacting to produce some interesting, complicated behaviour. The investigation of computer science and collective phenomena both involve the study of complicated behaviours emerging from simple systems following simple rules. This is particularly so in models of computation such as cellular automata or object oriented programming, in which the programs being executed do not have a natural sequential structure, but rather involve the parallel interaction of many relatively simple systems. The hope is that connections between the two fields can be found, based upon the analogy in the tasks the two fields attempt to accomplish. Indeed, some connections between the two fields are already known at the classical level (see for example [194] and references therein). However, little work investigating possible connections seems to have been done in the quantum case. We will return to this problem in Chapter 11 with some concrete proposals for investigation of the connections between these two areas.

I have repeatedly stressed the impact that physics has on the foundations of computer science, as it causes us to re-evaluate the fundamental models used in the study of information processing. Does the physics of information, especially quantum information, have a similar impact on fundamental physics?

The term *fundamental physics* itself has been the subject of considerable debate in recent years. On occasion, it appears merely to mean "my research is more important than yours". Two particularly strongly argued cases for what it means for a phenomenon to be fundamental have been presented by Anderson [4] and Weinberg [183].

Anderson's article, entitled "More is different", argues that essentially new principles appear at higher levels of complexity in physical systems, that cannot be deduced from the constituent parts alone. Anderson argues that the study of such phenomena is as fundamental as the study of particle physics or cosmology, which are traditionally regarded as the most fundamental parts of physics.

Weinberg takes a very different tack. He introduces what might be called "arrow diagrams" relating different realms of science to one another. There is an arrow from one field to another if the first field depends critically upon the second. For example, physical chemistry "points" to quantum mechanics, because the interactions of atoms and molecules are determined by the rules of quantum mechanics. Weinberg argues that fundamental phenomena are those which can not be reduced to

some simpler level; they have arrows pointing towards them, but none pointing away from them. Up to this point I'm with him, and believe his point of view dovetails nicely with Anderson's. Weinberg then goes on to assert, without any evidence that I can see, that particle physics and cosmology are the unique branches of science which have this property of irreducibility.

What seems to me to be going on here is a confusion of two separate issues. First is the question of whether or not a phenomenon is *universal*. Particle physics and cosmology study phenomena which are, without a doubt, universal. The second question is whether or not a phenomenon is *reducible* to a simpler theoretical level. For example, the energy levels of the Hydrogen atom can be explained quite well using simple quantum mechanics.

It seems to me that the term *fundamental* refers primarily to whether or not a phenomenon is reducible to some simpler level or not. If it is not, then our task as scientists must surely be to explain that phenomenon on its own terms.

It is instructive to consider the concrete example of thermodynamics. It is expected by many people that the principles of thermodynamics should be reducible to mechanics, yet despite decades of hard work such a reduction has never been generally achieved. It may be that such a reduction is in principle impossible. It is known, for example, that many behavioural properties of certain types of cellular automata can not be deduced merely by knowing their starting configurations and dynamics without performing a full simulation of the entire process [194].

What if such a situation were to obtain in the study of real phenomena: that the behaviour of those phenomena could not be deduced from their starting configurations and a detailed knowledge of their microscopic dynamics, by any means short of observing the actual ensuing dynamics? Would we give up attempts at explanation of those phenomena? Of course not! Our task then would be to discern higher level principles governing the behaviour of those systems, and to subject those principles to the same thorough empirical scrutiny which has been our wont at the microscopic level throughout the history of physics. This does not imply that we should give up the search for reductions of one theory to another, but rather, that we should acknowledge that such attempts may not always be successful, nor need they be possible, *even in principle*.

What then is the role of quantum information in fundamental science, especially fundamental physics? First, I believe it can be used to aid in the reduction of mesoscopic quantum phenomena to the level of elementary quantum mechanics. It is difficult to point to many situations where this has yet occurred, but I believe that is primarily because much of the field has been focused inward, on the development of basic tools. Recently there has been some indicators that this is occurring, such as the work of Huelga *et al* on using concepts from quantum information to develop better frequency standards [87].

Second, quantum information can directly assist the process of research into fundamental physics. One way of doing this is by throwing new light on old quantum principles which, as suggested earlier, is potentially a major stimulant of further progress in fundamental research. Another way in which quantum information can inform the progress of fundamental physics is to act as a source and catalyst for fundamental questions, such as the questions about the class of realizable measurements raised earlier in this Chapter, or to suggest new methods of approaching existing questions, such as Preskill's recent suggestion [143] that quantum error correction could provide a missing link between Hawking's claim [78] that at the fundamental level nature may be non-unitary, and the unitarity which appears to be the rule in all experimental work done to date. More precisely, Preskill has proposed that this apparent contradiction may be caused by some sort of "natural" quantum error correction, in which nature is non-unitary at very small length scales, but this non-unitarity gradually becomes less important at longer length scales.

Obviously, it is not possible to say with any degree of certainty how quantum information will affect fundamental physics in years to come. Yet I hope to have convinced you that quantum

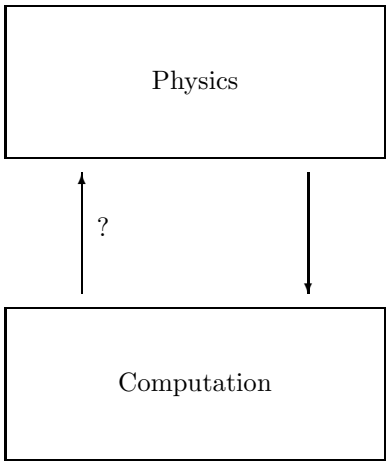

Figure 1.1: Adaptation of the meaning circuits proposed by Wheeler [186] and Landauer [107].

information is a subject worth thinking about in connection with fundamental physics, and that it has already resulted in some interesting work of a fundamental nature.

There is an interesting and related question one can ask about the comprehensibility of physical laws, which to my knowledge was first raised by Wheeler and Landauer, two of the first researchers to appreciate the deep connections between physics and computer science. In [186] and [107] they each proposed a "meaning circuit" to represent the connections between physics and computation. An adapted version of these circuits is shown in figure 1.1.

The circuit illustrates two connections between physics and computation. One is the observation that the laws of physics determine the scope of possible computational processes. This is an observation that we have discussed at length in this introductory Chapter, and is the founding insight for the entire Dissertation. I don't believe we can reject this part of the circuit without rejecting the founding principle of physics, namely that the world is essentially orderly, being governed by some set of laws.

The second part of the meaning circuit may be encapsulated in a question: *Are the consequences of the fundamental laws of physics computable?* The answer to this question depends on what is computable. That, as we have seen, depends on what the laws of physics are, so the question has an interesting self-referential nature. Another way of stating the question is: *Do the laws of physics allow the existence of structures capable of comprehending those laws?*

To make progress in physics, it is necessary to assume that the answer to this question is yes, at least in some limited domain. However, *a priori* there does not seem to be any especially good reason why the answer to this question ought to be yes, despite the empirical fact that a good deal about the universe *does* appear to be comprehensible. As Einstein noted, "*The most incomprehensible thing about the world is that it is comprehensible.*"

In my day to day work I, of course, assume that the laws of physics do allow the existence of structures capable of comprehending those laws; there wouldn't be much point to my work otherwise. Nevertheless, there are some interesting, amusing, and possibly even fruitful speculations one can engage in, by questioning this assumption.

One observation is that human beings may not naturally exploit the full information pro-

cessing power provided by the laws of physics. In particular, it appears as though information processing devices based upon quantum mechanics may be intrinsically much more powerful than devices which process information according to the principles of classical computer science. It is usually assumed, with considerable supporting evidence, that the human mind processes information in ways adequately modeled by classical models of information processing; this is is sometimes known as the *computational hypothesis* of cognitive science [130, 179]. Might it be possible that quantum mechanical devices for computation may exploit their additional processing power to achieve a more complete comprehension of the world than is possible using classical computational devices?

This line of speculation may be restated in the language of computer science. Roughly speaking, the class of problems known as NP in computer science are those for which a *solution* to the problem can be checked efficiently, that is, in polynomial time on a classical computer. For example, solutions to the well-known traveling salesman problem can be be checked quickly on a classical computer: given a potential solution, one simply checks whether that path length is less than the desired bound on path lengths. At the 1997 DIMACS Quantum Computing Tutorial and Workshop, Peter Shor commented on one possible relationship between classical and quantum complexity classes during his talk. Shor indicated a suspicion that there may be problems which can be solved efficiently on a quantum computer, for which even the solution may not be efficiently checkable on a classical computer. This would be truly remarkable if correct, and would open up the possibility of having a quantum oracle that can tell you what the answer to certain problems are, but not give you a proof of the answer which you can efficiently check. Might it be that the laws of physics can be comprehended by intelligences which do quantum information processing, yet some aspects of those laws remain uncomprehended by beings which only utilize classical models of computation?

Let us return from the far edges of speculation to the more practical concern of understanding the effect quantum information will have on technology.

In 1965, Gordon Moore presented a now-famous talk[5] in which he made a variety of predictions about how computer power would behave over the coming years. "Moore's Law" is quoted in a number of different forms; perhaps the most famous form is the economic form, that computer power will double for constant cost every two years or so[6]. We are more interested in the closely related physical forms of Moore's law, which state, for example, that the number of atoms needed to represent one bit of information should halve every two years or so. This prediction has, indeed, been borne out [50, 192] over the past thirty years.

Extrapolating this trend, it has been predicted [192] that around the year 2020, bits will be stored in individual atoms. At that level, we would certainly expect quantum mechanical effects to become very important. Techniques for quantum control[7], motivated in part by potential applications to quantum information processing, will be necessary to build components at that scale, even if quantum effects are not harnessed in the information processing model being implemented.

There is a second physical aspect to Moore's law. The amount of heat that may be dissipated by a chip with a given surface area per unit time is roughly a constant, without making use of elaborate refrigeration techniques. Thus, if the number of components being squeezed onto the chip is increasing, and the speed of logical operations on the chip is increasing, then the amount of heat that is being dissipated per logical operation must necessarily be decreasing. Once again,

[5] To my surprise, I have not been able to determine where this talk was given, or whether a written record of the talk exists.

[6] The quoted doubling time varies quite a bit depending from source to source, with the usual range of figures being one to two years.

[7] Depending on what one counts as quantum control, this is already a huge field. [122, 180, 119] are references on quantum control that are directly motivated by concerns of quantum information processing, and which provide an entry into the wider literature.

heat dissipation per logical operation has been following its own version of Moore's Law, halving approximately every twelve months [50, 192].

If current trends continue, by 2020, the amount of heat being dissipated per logical operation will be kT, which Landauer [106] has shown is the fundamental limit for irreversible logical operations. All modern computers are based upon such irreversible operations, so without radical changes in the architecture being used, we will hit a limit to computation set by heat dissipation requirements. Fortunately, a radical alternative does exist: reversible computation. Lecerf [110] and Bennett [14] have shown that it is possible to do universal computation when restricted to reversible logical operations, with negligible cost in terms of computational resources.

One way of accomplishing the big switch to reversible computation would be to move to quantum computation. Ideally, quantum computers are reversible devices at the atomic level. As we have seen, this is exactly what will be required for further progress, if current trends continue for another twenty years. A potential problem with this solution is that both quantum and classical reversible computers will require error correction techniques. As we discuss in Chapter 9, error correction itself necessarily dissipates energy, at a rate determined by the fundamental error rate in the information processing components. It seems likely that reversible classical information processing components will have a much lower fundamental error rate than their quantum components, which would make it much easier to switch to reversible computing rather than fully quantum computing. On the other hand, quantum computers have significant advantages over classical reversible computers, so it may make more sense to ameliorate the cost of switching to quantum computers by absorbing the funds that would have been necessary in any case to make the switch to classical reversible computation.

Until now, we have concentrated on the eventual *necessity* of taking quantum effects into account, if present trends in computer hardware are to continue. There is substantial economic incentive for the trends to continue, so it does not seem unreasonable to conclude that large semiconductor companies may eventually put serious effort into understanding and harnessing quantum effects.

Of course, the great promise of quantum information processing is to enable information processing tasks that are either intractably difficult, or downright impossible, using classical information processing techniques. At present, the most exciting potential applications known for quantum information are the ability to factor large composite numbers [165, 163], which in principle enables many currently popular cryptographic systems to be broken, and quantum cryptography [88, 16], which ironically enables unbreakable cryptographic systems. The attractiveness of both is derived from the widespread interest in private communications, for example, for financial transactions. I do not believe that either factoring or quantum cryptography is a truly "killer application" which makes the development of large scale quantum information processing imperative. One of the chief unknowns in the future of quantum information is whether such killer applications are possible, and if so, how they are to be found. This is a subject we will return to and discuss in more detail in Chapter 11.

What of the effect of quantum information on other fields? It is difficult to assess that effect until it becomes more clear what diagnostic use the tools of quantum information are. It is possible that techniques developed within the field of quantum information such as quantum error correction and quantum process tomography will provide much more precise information about the noise processes taking place at the atomic level than is currently known. If this hope is fulfilled, quantum information may have an enabling effect for fields such as biotechnology [73], quantum electronics [199], and molecular nanotechnology [62], whose eventual effectiveness depends, to some extent, on a detailed understanding of the processes taking place at the atomic level.

Let us conclude the Chapter by noting an amusing and optimistic "quantum corollary" to

Moore's law. In general, it is very difficult to simulate a quantum system on a classical computer. The difficulty of doing so rises exponentially with the number of qubits[8] in the system being simulated. In practice, what this means is that at least twice as much classical computer power is required to simulate a quantum system containing one more qubit. In fact, the problem may even be worse than that. The amount of memory required to store the quantum state at least doubles with each additional qubit, however the time required to do the simulation also rises because many more computational paths must be accounted for.

On a quantum computer, however, it is known how to efficiently simulate a wide class of quantum systems [117, 1, 200]. Very roughly, we may say that for this class of problems, quantum computers are keeping pace with classical computers provided a *single qubit* can be added to the quantum computer for every classical doubling period, according to Moore. This corollary should not be taken too seriously, as the exact nature of the gain, if any, of quantum computation over classical will not be clear until some major problems in computational complexity are resolved. Nevertheless, this is an amusing heuristic statement that helps convey why we should be interested in quantum computers, and hopeful that they will one day be able to outperform the most powerful classical computers, at least for some applications.

Summary of Chapter 1: The physics of information

- Information is physical. Each physical theory may be treated as the basis for a theory of information processing, with possible differences in resulting computational power.
- Quantum information is the study of the achievable limits to information processing possible within quantum mechanics.
- There is more than one type of information available to be processed in quantum mechanics, including *coherent quantum states, classical information*, and *quantum entanglement*.
- The promise of quantum information is to reveal new information processing capabilities beyond what is possible in traditional models of information processing, and to inform us as to the limits of quantum mechanics as a means for information processing.
- Quantum corollary to Moore's law: for certain applications, quantum computers need only increase in size by one qubit every two years, in order to keep pace with classical computers.

[8] The base two logarithm of the number of Hilbert space dimensions in non-qubit systems.

Chapter 2

Quantum information: fundamentals

This Chapter introduces many of the fundamental notions of quantum information, emphasizing notation, terminology, and simple examples. It is assumed that you are familiar with elementary quantum mechanics, in particular, the bra-ket notation for state vectors. The Chapter begins with an introduction to the fundamental unit of quantum information, the quantum bit, or *qubit*, and then gives two important examples of quantum information processing – superdense coding and quantum teleportation. With these concrete examples in hand, we introduce a general model of quantum information processing, the *quantum circuit*, or *quantum computing* model. This model is an attempt to formulate a general framework for the description of quantum information processing, and will be used as a tool in the description of many of the information processing tasks we discuss in the Dissertation. The Chapter concludes with an overview of the challenges facing an experimentalist wishing to do quantum information processing in the laboratory, a brief description of some of the technologies for quantum information processing which have been proposed or implemented, and a description of a new experimental implementation of quantum teleportation using liquid state nuclear magnetic resonance.

Before we begin the Chapter proper, a remark about notation. State vectors will be written in the standard bra-ket notation. Often, however, we will have occasion to use density operators. The standard notation for density operators is sometimes inappropriate when discussing composite systems. For example, a composite system consisting of two parts, A and B, will have three density operators associated to it and its various parts, ρ^A, ρ^B and ρ^{AB}. In addition, we will often wish to compare two or more different density operators on the same system, and may be interested in the states of the various system at different times. All this adds up to a mess of notation, with primes, subscripts and superscripts. For that reason, where it is clear, I often drop the ρ, and simply write A, B, and AB to indicate the density operators associated with the corresponding systems.

2.1 Quantum bits

The simplest quantum mechanical system has a two dimensional complex state space. Suppose we single out an orthonormal basis set in the state space of such a system, and label the basis vectors $|0\rangle$ and $|1\rangle$. Then an arbitrary pure state of the system has the form

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{2.1}$$

where α and β are complex numbers which must satisfy the condition $|\alpha|^2 + |\beta|^2 = 1$ in order for $|\psi\rangle$ to be correctly normalized.

This two dimensional quantum system is known as the *quantum bit* or *qubit* [154], by analogy with the *bit*, the fundamental unit of information in the classical theory of information processing. The states $|0\rangle$ and $|1\rangle$ are known as the *computational basis states*. They are merely reference states; it does not matter how they are chosen, just that we agree upon which states they are. In abstract discussions of quantum information processing, the computational basis states are no more than a fixed reference set of orthonormal basis states. In discussions of real physical systems implementing qubits, it is usual to pick the computational basis states so that they correspond to some other physically interesting pair of states. For example, in nuclear magnetic resonance implementations of a single qubit on a nucleus of spin 1/2, it is usual to identify the $|0\rangle$ and $|1\rangle$ states with the magnetic eigenstates of the spin corresponding to the large constant applied magnetic field.

It is instructive to compare bits and qubits. A bit can be in one of two states, 0 or 1. A qubit can be in a continuum of states, described by the complex numbers α and β. It is possible, in principle, to distinguish the 0 and 1 states of a bit. It is not possible, in general, to distinguish non-orthogonal states of a quantum system. For example, if we prepare a qubit in one of the two states $|0\rangle$ and $(|0\rangle + |1\rangle)/\sqrt{2}$, it can be shown that it is not possible to perform a measurement on that system which will reliably tell which of these two states was prepared[1].

By contrast, if we ensure that a qubit is always kept in the $|0\rangle$ or $|1\rangle$ state, then it is always possible to determine which state the system is in. Indeed, it turns out that all the information processing tasks which can be done with bits can also be done with qubits, provided the qubit remains in one of the two states $|0\rangle$ or $|1\rangle$. Thus, information processing models based upon the qubit are *at least* as powerful as models based upon the bit.

There are several items of terminology related to qubits which we ought to agree upon now. Four standard operators acting on a single qubit are the *Pauli sigma operators*, defined by

$$I \equiv \sigma_0 \equiv \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}; \quad X \equiv \sigma_1 \equiv \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}; \tag{2.2}$$

$$Y \equiv \sigma_2 \equiv \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}; \quad Z \equiv \sigma_3 \equiv \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \tag{2.3}$$

where these matrices are written in the computational basis $|0\rangle, |1\rangle$. The standard notation for the Pauli operators is σ_i; we will more often omit the redundant σ, and just write I, X, Y or Z instead.

The Pauli operators form a basis set for the vector space of operators on a single qubit. In particular, an arbitrary operator A acting on a single qubit can be written uniquely in the *Bloch representation*,

$$A = \sum_{i=0}^{3} a_i \sigma_i. \tag{2.4}$$

The Bloch representation has a particularly attractive form for density operators of a single qubit. Such a density operator can be written in the form

$$\rho = \frac{I + \vec{\lambda} \cdot \vec{\sigma}}{2}, \tag{2.5}$$

where $\vec{\lambda} = (\lambda_x, \lambda_y, \lambda_z)$ is the *Bloch vector* for the state, characterized by the requirement that the vector is real and satisfies $\|\vec{\lambda}\| \leq 1$.

[1] See section 6.1 for further discussion of this point.

The Pauli operators are our first examples of *quantum gates*. Classical information processing is accomplished by various *logic gates* which act on the bits being processed. Similarly, quantum information processing is accomplished by quantum gates. Quantum gates are operations acting on a fixed number of qubits. For example, the Pauli operators represent unitary evolutions which may take place on a single qubit. The X Pauli operator is often known as the *quantum not gate*, as it flips the computational basis states, $X|0\rangle = |1\rangle$ and $X|1\rangle = |0\rangle$, much as the classical not gate interchanges 0 and 1. The Z Pauli operator is often known as the *phase flip gate*, as it flips the relative phase of the computational basis states, $Z|0\rangle = |0\rangle$ and $Z|1\rangle = -|1\rangle$. At present there is no widely accepted term for the Y operator. I is, of course, the identity gate.

Two more quantum gates which are of great importance are the *Hadamard* and *phase shift* gates. These gates are defined, respectively, as follows:

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \tag{2.6}$$

$$S = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}. \tag{2.7}$$

Note that $S^4 = Z$, $HZH = X$, and $ZX = Y$, up to a global phase[2], so the Hadamard gate and phase shift together can be used to generate any of the Pauli operators. Later, we will introduce a two qubit gate, the controlled not gate. The controlled not gate, the Hadamard gate and the phase shift gate together form a *universal set* – any unitary operation can be approximated arbitrarily well making use of only these gates.

There are many reasons the qubit is regarded as the fundamental unit of quantum information. It is the simplest quantum mechanical system, and is quite easily analyzed. Moreover, the state space of any finite dimensional quantum system can be understood to be composed of a number of qubits. In this respect, the qubit closely resembles the classical bit. It is possible to formulate classical information processing in terms of *trits*, for example, which are classical systems taking the three values $0, 1$ and 2. In certain systems, it may even be more natural to do the analysis this way. However, little is lost from the theoretical point of view by regarding a trit as being composed of two bits, in which only the three states $00, 01$ and 10 are accessible. Similarly, a three dimensional quantum system can be regarded as essentially identical to a pair of qubits in which the state is guaranteed to be in the space spanned by the states $|00\rangle, |01\rangle$ and $|10\rangle$. For all these reasons, and others which will become apparent as we move deeper into quantum information, the qubit is regarded as the fundamental unit of quantum information.

2.2 Superdense coding

There is a simple but important example of quantum information processing known as *superdense coding* [19] which is explained in this section. This example shows that there are information processing tasks which can be performed with qubits which do not have natural analogues in terms of bits.

Superdense coding involves two parties, conventionally known as "Alice" and "Bob", who are a long way away from one another. Their goal is to transmit some classical information from Alice to Bob. Suppose Alice is in possession of two classical bits of information which she wishes to send Bob, but is only allowed to send a single qubit to Bob. Can she achieve her goal?

[2]More precisely, $ZX = iY$. In quantum mechanics, global phase factors such as the i can be ignored.

Superdense coding tells us that the answer to this question is yes. Suppose Alice and Bob initially share a pair of qubits in the entangled state

$$|\psi\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}}. \tag{2.8}$$

Alice is initially in possession of the first qubit, while Bob has possession of the second qubit. Note that this is a fixed state; there is no need for Alice to have sent Bob any qubits in order to prepare this state. Instead, some third party may prepare the entangled state ahead of time, sending one of the qubits to Alice, and the other to Bob.

By sending her single qubit to Bob, it turns out that Alice can communicate two bits of classical information to Bob. Here is the procedure she uses. If she wishes to send the bit string "00" to Bob then she does nothing at all to her state. If she wishes to send "01" then she applies the quantum not gate, X, to her qubit. If she wishes to send "10" then she applies the phase flip, Z, to her qubit. If she wishes to send "11" then she applies the iY gate to her qubit. The four resulting states are

$$00 : |\psi\rangle \rightarrow \frac{|00\rangle + |11\rangle}{\sqrt{2}} \tag{2.9}$$

$$01 : |\psi\rangle \rightarrow \frac{|10\rangle + |01\rangle}{\sqrt{2}} \tag{2.10}$$

$$10 : |\psi\rangle \rightarrow \frac{|00\rangle - |11\rangle}{\sqrt{2}} \tag{2.11}$$

$$11 : |\psi\rangle \rightarrow \frac{|01\rangle - |10\rangle}{\sqrt{2}}. \tag{2.12}$$

These four states are known as the *Bell basis*, after John Bell, who did so much to emphasize the importance of entanglement [12]. Notice that the Bell states form an orthonormal basis, and can therefore be distinguished by an appropriate quantum measurement. Alice now sends her qubit to Bob, giving Bob possession of both qubits. By doing a measurement in the Bell basis Bob can determine which of the four bit strings Alice sent.

This remarkable prediction of quantum mechanics has been given a partial experimental validation by Mattle *et al* using entangled photon pairs [124]. In the experiment, a trit of classical information was sent using photon polarization as the qubit. It was only possible to send a trit, rather than two bits, because with the measurement scheme used, the experimentalists were unable to distinguish between the states corresponding to 00 and 01, above.

It is surprising enough that a two level quantum system can be used to transmit two bits of classical information, however there is another remarkable aspect to this procedure. Suppose Alice sends her qubit to Bob, but the qubit is intercepted on the way by a third party, Eve. Examining the four states (2.9)–(2.12), we see that in each case, the reduced density operator associated with the first qubit is the same, the completely mixed state $I/2$. Because the reduced states are the same regardless of which state was prepared, Eve can infer nothing about the information Alice is trying to send by examining the qubit she has intercepted. The intercepted qubit contains essentially no classical information; rather, the classical information is contained jointly by the two qubits.

Superdense coding is an example of how quantum and classical information can be combined in an interesting way. In Chapter 6 we will return to study the limits to superdense coding in a much more detailed fashion, along the way to some results about the efficiency of distributed computations in quantum mechanics.

2.3 Quantum teleportation

The medium is the message.

– Marshall McLuhan [125]

Superdense coding shows that quantum information may be used in an interesting way as the medium for transmission of classical information. An even more remarkable effect, quantum teleportation [18], shows that classical information and entanglement can be used as the medium for transmission of a quantum state.

Suppose Alice is living in London and wishes to send a single qubit to Bob, who is living in New York. There are many different ways Alice could do this. One method is for Alice to send a description of her state to Bob, who can then create that state in New York. This method has two major disadvantages. First, quantum states are specified using sets of complex numbers. For quantum systems of many qubits it requires a huge number of classical bits to specify the state to reasonable accuracy. The cost of transmitting these classical bits may be considerable. Second, suppose the state of Alice's system is not known to Alice. The situation then is even worse, because it is not possible, even in principle, for Alice to determine the state of her system. There is no way she can send her system to Bob by sending Bob a classical description.

A second method is to physically move the quantum system from London to New York. For example, a photon could be sent down a highly idealized fiber optic from London to New York. This method also suffers from two major disadvantages. First, it may simply be very difficult to reliably send qubits from London to New York. The channel used to do so may degrade over time, or it might be unreliable to begin with. Second, if the qubit being sent was carrying information that information could be intercepted by a malevolent third party.

Quantum teleportation is a method for moving quantum states from one location to another which suffers from none of these problems. Suppose Alice and Bob share a pair of qubits which are initially in the entangled state $(|00\rangle + |11\rangle)/\sqrt{2}$. In addition, Alice has a system which is in some potentially unknown state $|\psi\rangle$. The total state of the system is therefore

$$|\psi\rangle \left(\frac{|00\rangle + |11\rangle}{\sqrt{2}} \right). \tag{2.13}$$

By writing the state $|\psi\rangle$ as $\alpha|0\rangle + \beta|1\rangle$ and doing some simple algebra, we see that the initial state can be rewritten as

$$(|00\rangle + |11\rangle)|\psi\rangle + (|00\rangle - |11\rangle)Z|\psi\rangle + (|01\rangle + |10\rangle)X|\psi\rangle + (|01\rangle - |10\rangle)XZ|\psi\rangle. \tag{2.14}$$

Here and throughout the remainder of this section we omit normalization factors from the description of quantum states.

Suppose Alice performs a measurement on the two qubits in her possession, in the Bell basis, consisting of the four orthogonal vectors, $|00\rangle + |11\rangle, |00\rangle - |11\rangle, |01\rangle + |10\rangle, |01\rangle - |10\rangle$, with corresponding measurement outcomes which we label $00, 01, 10$ and 11. From the previous equation, we see that Bob's state, conditioned on the respective measurement outcomes, is given by

$$00 : |\psi\rangle; \quad 01 : X|\psi\rangle; \quad 10 : Z|\psi\rangle; \quad 11 : XZ|\psi\rangle. \tag{2.15}$$

Therefore, if Alice transmits the two classical bits of information she obtains from the measurement to Bob, it is possible for Bob to recover the original state $|\psi\rangle$ by applying unitary operators inverse to the identity, X, Z and XZ, respectively. More explicitly, if Bob receives 00, he knows his

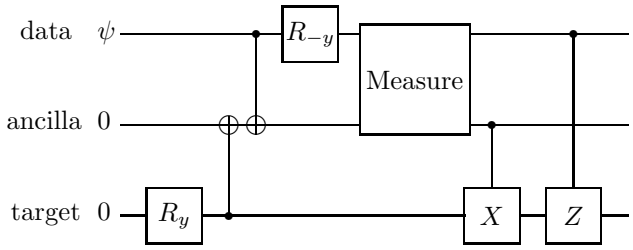

Figure 2.1: Circuit for quantum teleportation. The measurement is in the computational basis, leaving the measurement result stored in the data and ancilla qubits. R_y and R_{-y} denote rotations of 90 degrees about the y and $-y$ axes on the Bloch sphere.

state is $|\psi\rangle$, if he receives 01 then applying an X gate will cause him to recover $|\psi\rangle$, if he receives 10 then applying a Z gate will cause him to recover $|\psi\rangle$, and if he receives 11 then applying an X gate followed by a Z gate will enable him to recover $|\psi\rangle$.

This completes the teleportation process.

It is interesting to note that teleportation involves the transmission of only *two bits* of classical information. This is despite the fact that in general it takes an infinite amount of classical information to describe the state to be teleported. Furthermore, the success of the teleportation procedure did not in any way depend upon Alice knowing anything about the quantum state she was sending. Even more remarkably, we see from equation (2.14) that each of the four Bell states appears with equal weight in the superposition making up the initial state. Thus, the four measurement results have equal probabilities 1/4, *independent of the initial state* $|\psi\rangle$. Because the probability is independent of the state $|\psi\rangle$, neither Alice nor anybody else can infer anything about the identity of the state being teleported from the measurement outcome.

Quantum teleportation can be recast in the language of quantum gates which we met briefly earlier in this Chapter. A *quantum circuit* implementing teleportation is shown in figure 2.1 [30]. The three lines traversing the circuit from left to right represent the three qubits involved in teleportation. The top line represents the initial state which Alice wishes to teleport. We shall refer to it as the *data qubit*. The second line represents the qubit which Alice uses to share the initial entanglement with Bob, which we shall call the *ancilla qubit*. The third line represents Bob's qubit, which we shall call the *target qubit*.

The quantum circuit is read from left to right. The input state to the quantum circuit is assumed to be the product state $|\psi\rangle|00\rangle$, and the first two gates in the circuit are used to create the entanglement between Alice and Bob. The very first gate is a 90 degree rotation about the y axis on the Bloch sphere[3]. This takes the state $|0\rangle$ to the state $(|0\rangle + |1\rangle)/\sqrt{2}$. The second gate is known as the controlled not gate. As can be seen from the figure, the controlled not gate involves two qubits, which we shall refer to as the *control* and *data*[4] qubits. In this particular gate, the control qubit is on the bottom line, and the data qubit is on the second line.

The controlled not gate is a unitary gate whose action is to flip the data qubit if the control

[3]In general, a rotation by θ degrees about the $\vec{n}$ axis on the Bloch sphere is defined to be $\exp(-i\theta\vec{n}\cdot\vec{\sigma}/2)$, where $\vec{n}$ is a unit vector, in this case $(0,1,0)$, and $\vec{\sigma} = (X,Y,Z)$ is a vector whose entries are the Pauli sigma operators. Therefore, a 90 degree rotation about the y axis corresponds to the operator $(I - iY)/\sqrt{2}$.

[4]The data qubit for a controlled not gate is sometimes known as the target qubit, not to be confused with the target qubit of our circuit!

qubit is set to $|1\rangle$, and to leave the data qubit alone if the control qubit is set to $|0\rangle$. Symbolically,

$$|a\rangle|b\rangle \stackrel{\text{c-not}}{\longrightarrow} |a\rangle X^a|b\rangle, \tag{2.16}$$

where a and b are 0 or 1. This defines the action of the controlled not on a basis, and thus on all states.

After the controlled not is applied, the state of the system is seen to be

$$|\psi\rangle \frac{|00\rangle + |11\rangle}{\sqrt{2}}. \tag{2.17}$$

That is, the first two gates create the necessary entanglement between the target and ancilla qubits.

The next step of teleportation is to perform a measurement on the data and ancilla qubits in the Bell basis. The way this is accomplished is to do two gates which rotate the Bell states into the computational basis, and then to perform a measurement in the computational basis. The rotation is accomplished by performing a controlled not from the data qubit to the ancilla qubit, followed by a rotation by -90 degrees about the y axis on the data qubit. The effect of these transformations is as follows:

$$\frac{|00\rangle + |11\rangle}{\sqrt{2}} \rightarrow \frac{|00\rangle + |10\rangle}{\sqrt{2}} \quad \rightarrow |00\rangle \tag{2.18}$$

$$\frac{|00\rangle - |11\rangle}{\sqrt{2}} \rightarrow \frac{|00\rangle - |10\rangle}{\sqrt{2}} \quad \rightarrow -|10\rangle \tag{2.19}$$

$$\frac{|01\rangle + |10\rangle}{\sqrt{2}} \rightarrow \frac{|01\rangle + |11\rangle}{\sqrt{2}} \quad \rightarrow |01\rangle \tag{2.20}$$

$$\frac{|01\rangle - |10\rangle}{\sqrt{2}} \rightarrow \frac{|01\rangle - |11\rangle}{\sqrt{2}} \quad \rightarrow -|11\rangle. \tag{2.21}$$

Thus, a measurement in the computational basis will give the result $00, 01, 10$, or 11, corresponding to one of the four Bell states. Moreover, the state of the data and ancilla qubits after the measurement is a computational basis state whose value records the result of the measurement; that is, which Bell state was measured.

Making use of this fact, we apply operations to the target qubit, conditional upon the measurement result. First, we apply an X gate to the target qubit, conditional on the ancilla qubit being set, then a Z gate to the target qubit, conditional on the data qubit being set. Thus, the four possible outcomes are: if 00 is measured, then the identity transformation is applied to the target; if 01 is measured then X is applied; if 10 is measured then Z is applied to the target; and if 11 is measured then ZX is applied to the target. Notice that this sequence of operations corresponds exactly to the sequence of operations necessary for quantum teleportation. It is interesting to note further that the measurement step can be removed from the circuit and the state of the data qubit will still be transferred to the target qubit. However, this is much less impressive than the full teleportation operation, in which our intuition incorrectly tells us that the quantum state initially on the data qubit is irreversibly destroyed by the measurement process. This completes our description of quantum teleportation in the language of quantum circuits.

Quantum teleportation is an important elementary demonstration of quantum information theory. Later in this Chapter we discuss the experimental implementation of quantum teleportation, and in Chapter 3 we return to look at quantum teleportation from a completely different angle, as a noisy quantum channel. Finally, in Part II of the Dissertation we will repeatedly use teleportation as an elementary operation as part of more sophisticated quantum information processing operations. These and many other uses emphasize the role quantum teleportation has as an exemplar useful for the study of more complex forms of quantum information processing.

2.4 Quantum computation

The theory of quantum computation is an attempt to capture the essential elements of a theory of quantum information processing in a single unified theory. I say "attempt" because it is not yet clear that the theory of quantum computation provides a complete account of the information processing capabilities afforded by quantum mechanics.

This section describes a single model of quantum computation, the *quantum circuit* model. Other, equivalent, formulations of quantum computation have also been proposed, but these will not be discussed in any detail here. Without further ado, here is an outline of the quantum circuit model of quantum computation:

1. **Classical resources:** The quantum computer consists of two parts, a classical part and a quantum part. In principle, there is no need for the classical part of the computer, but in practice certain tasks may be made much easier if parts of the computation can be done classically. For example, many schemes for quantum error correction are likely to involve classical computations in order to maximize efficiency. While classical computations can always be done, in principle, on a quantum computer, it may be more convenient to perform the calculations on a classical computer.

2. **A suitable state space:** We assume that the quantum part of the computer consists of some number, n, of qubits. The state space is thus a 2^n dimensional complex Hilbert space. Product states of the form $|x_1, \ldots, x_n\rangle$, where $x_i = 0, 1$, are known as *computational basis states* of the computer. We sometimes write $|x\rangle$ for a computational basis state, where x is the number whose binary representation is $x_1 \ldots x_n$.

3. **Ability to prepare states in the computational basis:** It is assumed that any computational basis state $|x_1, \ldots, x_n\rangle$ can be prepared in at most n steps.

4. **Ability to perform quantum gates:** It is assumed that it is possible to perform the *Hadamard gate* and the $\pi/4$ phase shift gate on any single qubit of the quantum computer. It is assumed that it is possible to perform the *controlled not* gate on any pair of qubits in the quantum computer. Recall that these gates are defined in the computational basis as follows:

   - The Hadamard gate:
   $$H \equiv \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \tag{2.22}$$
   - The $\pi/4$ phase shift gate:
   $$S \equiv \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}. \tag{2.23}$$
   - The controlled not gate:
   $$C \equiv \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \tag{2.24}$$

5. **Ability to perform measurements in the computational basis:** Measurements may be performed in the computational basis of one or more of the qubits in the computer.

6. **Algorithm to build the quantum circuit:** Suppose we wish to solve a problem using the quantum circuit model of computation. As an example, suppose we wish to factor numbers in the quantum circuit model. Then given an input for the problem, in this case the number to be factored, there must be a procedure telling us how to build the quantum circuit to perform the desired computation. That is, we must have an algorithm which, given the input, describes how many qubits will be needed to do the computation, which computational basis state must be prepared, what gates must be applied during the computation, and when those gates are to be applied, what measurements are to be performed during the computation, and a specification of what measurement results are to be regarded as output from the computation. This requirement – that the structure of the quantum circuit be specified by a classical algorithm – is known as the *uniformity requirement* for quantum computation. Without imposing this important requirement, many impossible tasks would become trivial within the quantum circuit model, or even in the classical circuit model of computation [139].

This model of computation is equivalent to many other models of computation which have been proposed, in the sense that other models result in similar resource requirements for the same problems. For example, one might wonder whether moving to a design based on three level quantum systems, rather than the two level qubits, would confer any computational advantage. Of course, although there may be some slight advantage in using three level quantum systems over two level systems, any difference will be essentially negligible from the theoretical point of view. At a less trivial level, the "quantum Turing machine" model of computation, a quantum generalization of the classical Turing machine model, has been shown to be equivalent to the model based upon quantum circuits [198, 25].

In what ways may the quantum circuit model of computation be criticized? How might it be modified? Perhaps my sharpest criticism of the quantum circuit model is that its basis, although expressed in terms of quantum mechanics, is not yet wholly rooted in fundamental physical law. The basic assumptions underlying the model are *ad hoc*, and do not seem to have been analyzed in the literature with respect to fundamental physical law, at least not in any great depth.

For example, it is by no means clear that the basic assumptions underlying the state space and starting conditions in the quantum circuit model are justified. Everything is phrased in terms of finite dimensional state spaces. Might there be anything to be gained by using systems whose state space is infinite dimensional? What about the assumption that the starting state of the computer is a computational basis state? We know that many systems in nature "prefer" to sit in highly entangled states of many systems; might it be possible to exploit this preference to obtain extra computational power? It might be that having access to certain states allows particular computations to be done much more easily than if we are constrained to start in the computational basis. Likewise, if measurements could be performed outside the computational basis, it might be possible to harness those measurements to perform tasks intractable within the quantum circuit model.

It is not my purpose here to do a detailed examination of the physics underlying the models used for quantum computation, although I believe that this is a problem well worth considerable time and effort. I wish merely to raise in your mind the question of the completeness of the quantum circuit model, and re-emphasize the fundamental point that information is physical, and in our attempts to formulate models for information processing we should always attempt to go back to fundamental physical laws. A very desirable goal for the future is to use fundamental physics to demonstrate or refute the following modern version of the Church-Turing thesis (see also [58]):

Any physically reasonable model of computation can be simulated in the quantum circuit model with at most polynomial overhead in physical resources.

At a more practical level, the quantum circuit model will be used in this Dissertation to

provide a language for the description of quantum information processing tasks. We have already done that, for example, in the description of quantum teleportation. It is only in the next section that the question of the completeness of the quantum circuit model of quantum information processing will be an important issue.

2.5 What quantum measurements may be realized?

Let us return to the question asked in section 1.2: What observables may be realized as measurements on a quantum system? In this section we discuss this problem from a somewhat different point of view than was done earlier. The point of view we take is that the quantum circuit model provides an essentially complete account of the information processing tasks, including measurement, that can be accomplished within quantum theory. As noted in the previous section, that this is a valid assumption has not yet been established beyond doubt, however, it will allow us to make progress on the question of determining what measurements may be performed within quantum mechanics.

We begin with the halting problem, and a proof that the halting problem is algorithmically unsolvable. The discussion is a little different to Turing's [177], since we allow probabilistic algorithms, which generalize the deterministic algorithms considered by Turing. The central outcome of our discussion is the same as Turing's though: the halting problem may not be solved by any algorithm, even a probabilistic algorithm.

Turing's key insight was to formalize what he meant by an algorithm. Essentially, Turing invented the modern concept of a programming language for his computers. An algorithm to compute a function is expressed in terms of a program, which takes as input a number, and outputs a number – the input and output of the function computed by the program. Strictly speaking, the functions computed by programs are *partial* functions, since it is possible that for some inputs a program will fail to ever halt; the function computed by the program is therefore undefined for that input.

A key point made by Turing is that his programs can be numbered $0, 1, 2, \ldots$. There is no need to explicitly give the details of Turing's model of computation here; it is well covered in computer science texts such as [57] and [49]. We need only imagine a computer running a program in a familiar language such as C or PASCAL, with the caveat that programs running on the computer, while finite, may make use of an arbitrarily large amount of scratch memory while running. We consider a slight generalization of Turing's model in that we assume that the computer is equipped with a good random number generator which generates either a zero or one with equal probability. This random number generator can be called as part of the algorithm. It is not difficult to see that even in this slightly generalized model of computation, it is still possible to number the possible programs, $0, 1, 2, \ldots$.

We define the halting function in the probabilistic model of computation as

$$h(x) \equiv \begin{cases} 0 & \text{if program } x \text{ halts with probability } < \frac{1}{2} \text{ on input } x \\ 1 & \text{if program } x \text{ halts with probability } \geq \frac{1}{2} \text{ on input } x. \end{cases} \tag{2.25}$$

Is there an algorithm which computes the halting function? More precisely, does there exist an algorithm which can compute the halting function better than just randomly guessing, that is, with probability of correctness greater than one half for each input? We will give a proof by contradiction that such an algorithm cannot exist, by assuming that such an algorithm does exist, and showing that it leads to a contradiction. More formally, for each possible input x, the algorithm, which we shall call HALT, outputs $h(x)$ with probability greater than one half. We make use of the algorithm for HALT to construct another program, which we call TURING, which calls HALT as a subroutine. In pseudocode:

```
program TURING(x)
y = HALT(x)
if y = 0 then
   return x and halt
else
   loop forever
```

We have assumed that HALT is computable, and thus the program TURING is also computable, and must have an associated program number, say t. What is the value of $h(t)$? Notice that $h(t) = 1$ if and only if TURING halts on input of t with probability at least one half. Inspecting the program for TURING, we see that this is true if and only if HALT$(t) = 0$ with probability at least one half, and thus, strictly greater than one half, by assumption. Finally, by definition this last is true if and only if $h(t) = 0$. That is, $h(t) = 1$ if and only if $h(t) = 0$, clearly a contradiction. Thus, our original assumption must have been wrong: there is no algorithm which can compute the halting function with success probability greater than one half for all inputs.

Having demonstrated that there is no algorithm capable of computing the halting function, even probabilistically, let us return to the problem of measurements in quantum mechanics. Suppose that, rather than wishing to implement a specific type of measurement, we wish to implement a *family* of measurements. We define a family of measurements to consist of a sequence $\mathcal{M} = \{M_1, M_2, \ldots\}$ of observables, where M_n is an observable on n qubits. Given n, is it possible, in general, to perform a measurement of M_n on n qubits?

The answer is no. In particular, define the *halting family* $\mathcal{M}$ of observables by

$$M_n \equiv \sum_{x=0}^{2^n-1} h(x)|x\rangle\langle x|. \tag{2.26}$$

The halting family can not be measured, even approximately, within the quantum circuit model of computation. To see why not, suppose that it is possible to measure this family of observables within the quantum circuit model. We will outline an algorithm for a Turing machine that will compute the halting function. The algorithm is very simple: given x, it chooses n greater than $\log x$, and then simulates the quantum circuit used to measure the observable M_n on n qubits, with the starting state for the circuit chosen to be $|x\rangle$. By carrying out the simulation to a high enough level of accuracy, we can ensure that the value of $h(x)$ can be read out from the output of the simulation.

We are left to conclude that it is *not* possible, in principle, to measure the halting family of observables within the quantum circuit model. If we assume that the quantum circuit model provides a complete description of the class of information processing tasks which may be performed in quantum mechanics then we are left to conclude that physical law does not allow measurement of the halting family of observables.

It is intriguing to consider the consequences if it were possible to measure the halting observable or, which is effectively the same thing, the halting family of observables. Perhaps there really exist in nature quantum processes which can be used to compute functions which are classically non-computable. It is far-fetched, but not logically inconsistent, to imagine some type of experiment - perhaps a scattering experiment - which can be used to evaluate the halting function.

Recognizing such a process poses some problems. How could we verify that a process computes the halting function (or any other non-computable function)? Because of the algorithmic unsolvability of the halting problem, it is not possible to verify directly that the candidate "halting process" does, in fact, computing the halting function. Nevertheless, one can imagine inductively verifying that the process computes the halting function. In principle, one could do this by running

a large number of programs on a computer for a long time, and checking that all the programs which halt are predicted to halt by the candidate halting process, and that programs predicted not to halt by the candidate halting process have not halted. Given sufficient empirical evidence of this sort, one could then *postulate* as a new physical law that the process computes the halting function.

Physically, the most important conclusion we can draw from this discussion is that there may be significant limitations on the class of observables which can be realized in quantum mechanics. Related restrictions apply for unitary dynamics [132]. These limitations may go considerably beyond the familiar limits of the type discovered by Heisenberg, although it is not yet clear precisely what class of measurements is realizable in quantum mechanics. In what future directions may this line of thinking be taken? The most obvious is to clarify the extent to which the quantum circuit model of computation is a complete framework for the description of quantum information processing. I believe this would be a long and difficult task, but well worth doing. If such a result could be established, then one could develop a theory of realizable measurements, along lines similar to recursive function theory in computer science [57, 139].

That concludes our discussion of realizable quantum measurements. In many ways it is a digression from the main stream of the Dissertation, but it is a digression that reinforces many of the points made in the main stream, and alerts us to some open problems in fundamental physics that I would very much like to see solved. Let us now turn to the more immediately practical topic of experimental quantum information processing.

2.6 Experimental quantum information processing

The theory of quantum information processing has progressed very quickly over the past twenty years. By contrast, experimental progress has been much slower, despite much ingenuity and effort on the part of experimentalists.

This section reviews the requirements that must be met in order to do interesting quantum information processing tasks, and describes in some detail one of the specific technologies proposed to perform quantum information processing. The section begins with a discussion of some of the general principles to be met by quantum information processors[5]. We then discuss in some detail the approach to quantum computing based upon liquid state nuclear magnetic resonance. The section concludes with an account of the use of nuclear magnetic resonance to accomplish quantum teleportation.

The specific requirements which must be met by a system which is to do quantum information processing depend upon the task which the system is to perform. For example, tasks such as superdense coding require a high level of control over single qubits, but only a small number of qubits in order to be accomplished. Optical methods have been used to successfully implement an impressive variety of quantum information processing tasks of this high precision-small size type, including quantum cryptography ([88], and references therein), a variant of superdense coding [124], and quantum teleportation [28, 29]. Given this impressive progress, it seems likely that optical methods will remain important for quantum information processing, at least in the short term.

These same optical methods are of little use in their present form for more general quantum information processing tasks. Purely optical methods do not appear to scale very well, and with present techniques it is very difficult to implement the non-linear optical interactions which are necessary for quantum logic. Physically, in order to achieve the interactions between photons neces-

[5] We follow Steane [171] in using the general term "quantum information processor" to describe any system that can be used to do quantum information processing, from the most elementary tasks, up to full-fledged quantum computation.

sary for quantum logic, there must be some other medium present to mediate the interaction, and presently known mediums for this interaction are not especially efficient. Moreover, it is difficult to store photons in a controlled fashion for long periods of time. These problems make it seem unlikely that photons will be the primary basis for large scale quantum information processors. In the near term, optical methods are likely to remain an important means for doing small scale investigations of quantum information processing, and it certainly seems reasonably likely that optical methods will have *some* role to play in other technologies for quantum information processing.

What general requirements are desirable in a system which is to be used for large scale quantum information processing? Obviously, the requirements to be met depend on the exact model of quantum information processing which is to be implemented; this is one of the reasons it is interesting to formulate different but equivalent models of quantum information processing. If the goal is to implement the quantum circuit model described in the previous section, then the following requirements must be met:

- The system must have a suitable n qubit state space.
- Ability to prepare the system in computational basis states.
- Ability to perform an appropriate universal set of gates on the system, for example, the controlled not, phase shift, and Hadamard gates.
- Ability to perform measurements in the computational basis.
- Precise external control over the system, allowing an arbitrary sequence of gates and computational basis state measurements to be performed on the system.

Finally, there is one additional requirement not directly related to the abstract theoretical model for the quantum circuit model, but of overwhelming practical importance: the ability to cope with noise. The performance of each of the above tasks will inevitably be imperfect, and quantum computers must be resilient in the face of such noise. In particular, the timescale t_c over which the coherent dynamics of the system takes place (roughly, the longest time required to perform one of the fundamental logical operations), must be very short compared to the timescale t_n over which the system's state is effectively messed up due to the effects of noise. Roughly speaking, the number of operations which can be done before a quantum computer becomes useless as a quantum computer is t_n/t_c. Thus, the goal is to find systems which maximize t_n while minimizing the time required for dynamics. In Chapter 9 we will investigate quantum error correcting codes which, it has recently been shown, can be used to effectively increase t_n/t_c for a quantum system, for little cost in the time required to do the coherent dynamics.

2.6.1 Proposals for quantum information processing

Many proposals have been made for systems capable of functioning as quantum information processors. Two of these proposals stand out as they have led to the successful implementation of simple quantum logical operations, and promise substantially more in the relatively near future. These proposals are based on the linear ion trap, originally proposed by Cirac and Zoller [42] and further developed by several groups of researchers [171, 127, 142], and the liquid state nuclear magnetic resonance (NMR) approach to quantum information processing. In this subsection we focus on a description of the NMR approach, in preparation for the next subsection, which describes the results of a collaboration with Knill and Laflamme to do quantum teleportation in NMR. It is also worth noting that a third technology, cavity QED [176], has been used to implement simple quantum logic.

This technology will not be reviewed here as the task of using this implementation of quantum logic to do more complex operations is even more formidable than for the ion trap, or NMR.

Methods for doing quantum information processing using liquid state NMR were proposed independently at about the same time by Cory, Fahmy and Havel [51], and by Gershenfeld and Chuang [69]. The scheme has since been applied to do numerous interesting quantum information processing tasks [38, 41, 53, 52, 89, 90, 105, 137].

The NMR method is unusual in that it makes use of a model of quantum information processing that is significantly different to the quantum circuit model of quantum computation. In particular, the computation is done in a bulk system, at room temperature. Therefore, the initial state of the system is not a pure computational basis state, but rather is a thermal mixture of states of the system. Furthermore, because of the bulk nature of the system it is not possible to do projective measurements on a single system, but rather, only ensemble averaged measurements can be made. Fortunately, both these problems can be circumvented in our effort to do quantum information processing.

The liquid state NMR approach to quantum information processing makes use of a large number of molecules dissolved in a solvent such as chloroform. For example, in experiments done in collaboration with Knill and Laflamme [137] at the Los Alamos National Laboratory, the molecule trichloroethylene, or TCE, was used. The structure of the molecule is shown in figure 2.2.

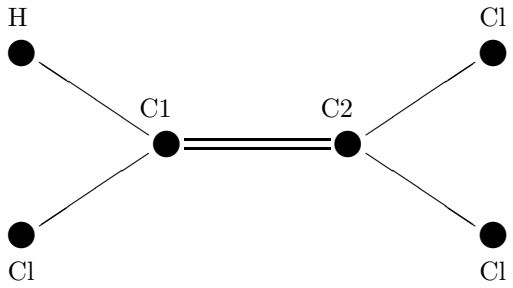

Figure 2.2: Schematic representation of the labeled TCE molecule. The two carbon atoms, C1 and C2, are ^{13}C isotopes, which a net nuclear spin of 1/2.

This molecule consists of two Carbon atoms, double chemically bonded, a Hydrogen atom, and three Chlorine atoms. The molecules are prepared in such a way that the Carbon atoms are actually the ^{13}C isotope, in order to give us a usable net spin 1/2 contribution from the nucleus. In our setup, the Chlorines are not usable because of the lack of a suitable detector.

The sample is placed in a large, homogeneous, static magnetic field, oriented in what we shall call the z direction. The field is as large as can be made with current technology for reasonable cost, typically in the range of 10 Tesla or so; in our experiments an 11.5 Tesla magnetic field was used. In the liquid state, the molecules in the sample tumble rapidly around, leading to a situation where interactions between the molecules can be ignored.

In this limit, the Hamiltonian describing the behaviour of the system is ($\hbar = 1$)

$$H = \sum_i \omega_i Z_i + \sum_{ij} J_{ij} Z_i Z_j, \tag{2.27}$$

where the second sum is over all pairs of spins. The values of the frequencies ω_i and J_{ij} depend on the particular spins being used; typically, $\omega_i \approx 10^8$-10^9 Hz, while for neighbouring spins $J_{ij} \approx 10^2$-10^3 Hz.

In the TCE molecule, the frequencies are as follows:

$$\omega_H \approx 500.133491\text{MHz}; \qquad \omega_{C1} \approx 125.772580\text{MHz}; \qquad \omega_{C2} \approx \omega_{C1} - 911\text{Hz} \tag{2.28}$$

$$J_{H\,C1} \approx 201\text{Hz}; \qquad J_{C1\,C2} \approx 103\text{Hz}. \tag{2.29}$$

The coupling frequencies between H and C2, as well as the Chlorines to H, C1 or C2, are much lower, on the order of ten Hertz for the former, and less than a Hertz for the latter. These couplings can be effectively removed by a technique known as *refocusing*, described below, and will be ignored in what follows. We will also ignore the Chlorines, as they were not visible in our experiment. Note that the frequencies of C1 and C2 are not identical; they have slightly different frequencies, due to the different chemical environments of the two atoms. This effect is known as the *chemical shift*.

In addition to the uniform magnetic field, it is possible to apply radio frequency (rf) pulses on resonance to each of the spins in a direction transverse to the direction of the uniform magnetic field, that is, in the $x - y$ plane. In a frame rotating about the z axis with the spins at respective frequencies ω_i, the Hamiltonian for this system can therefore be approximated as

$$H = \sum_{ij} J_{ij} Z_i Z_j + \sum_i P_{i\theta}(t)(\cos\theta\, X_i + \sin\theta\, Y_i), \tag{2.30}$$

where θ is some phase that may be externally controlled, and the strength $P_{i\theta}(t)$ of the rf pulse applied to spin i may also be controlled externally.

Using these external rf fields it is possible to perform single qubit rotations on individual nuclei in the molecule, by applying a field tuned to the appropriate resonance frequency. The necessary interactions happen fast enough that the contribution from the ZZ coupling between spins may be neglected. For our present purposes, it is sufficient to consider $\pi/2$ and π rotations about the $x, y, -x$ and $-y$ axes. For example, a $\pi/2$ rotation about the x axis has the effect

$$\exp(-i\pi X/4) = \frac{I - iX}{\sqrt{2}}. \tag{2.31}$$

A $-\pi/2$ rotation about the y axis has the effect

$$\exp(+i\pi Y/4) = \frac{I + iY}{\sqrt{2}}. \tag{2.32}$$

A π rotation about the x axis has the effect $\exp(-i\pi X/2) = -iX$. Similar observations may be made about the other possible rotations.

In the absence of externally applied rf fields, in the rotating frame the spins evolve according to the Hamiltonian

$$H = \sum_{ij} J_{ij} Z_i Z_j, \tag{2.33}$$

where the sum is over all pairs of interacting spins, (i, j). In many situations, it is desirable to be able to "turn off" one or more of these interactions. A clever technique known as *refocusing* allows this to be done. For simplicity, we explain how refocusing works in the specific case of the TCE

molecule. Suppose we wished to obtain a $\pi/2$ ZZ coupling between C1 and C2, with no interaction occurring between C1 and the Hydrogen.

In order to achieve, this, let t be any length of time. For example, we might choose t equal to the length of time necessary for a $\pi/2$ coupling between C1 and C2. Suppose we cause the following sequence of operations to occur:

1. Let the system freely evolve for a time $t/2$.
2. Apply a π pulse to H, in the X phase.
3. Let the system freely evolve for a time $t/2$.

The total evolution is thus

$$\begin{aligned} &\exp(-i(t/2)J_{C1\,C2}Z_{C1}Z_{C2})\exp(-i(t/2)J_{H\,C1}Z_HZ_{C1})(iX_H)\times \\ &\exp(-i(t/2)J_{H\,C1}Z_HZ_{C1})\exp(-i(t/2)J_{C1\,C2}Z_{C1}Z_{C2}). \end{aligned} \tag{2.34}$$

By the anticommutation relation $XZ + ZX = 0$ we see that

$$\begin{aligned} \exp(-i(t/2)J_{H\,C1}Z_HZ_{C1})(iX_H)\exp(-i(t/2)J_{H\,C1}Z_HZ_{C1}) &= \\ (iX_H)\exp(+i(t/2)J_{H\,C1}Z_HZ_{C1})\exp(-i(t/2)J_{H\,C1}Z_HZ_{C1}) &= iX_H. \end{aligned} \tag{2.35}$$

Thus, the total effect of this sequence of operations is

$$iX_H\exp(-itJ_{C1C2}Z_{C1}Z_{C2}). \tag{2.36}$$

That is, it is effectively as if a ZZ coupling of time t between C1 and C2 had occurred, together with a π rotation on the Hydrogen. The interaction between H and C1 has vanished; we say that it has been refocused.

We will use single qubit rotations and spin-spin couplings to perform unitary dynamics on our nuclear spins. Whether this forms a universal set for quantum computation depends upon the details of the molecule being considered; see [69] for a discussion of this point. For our much less grandiose purpose of doing quantum teleportation the interactions available are certainly sufficient to implement the quantum circuit for teleportation. The chief difficulty is perhaps that pulses applied to the two carbon nuclei are applied non-selectively. However, standard tricks based upon the chemical shift can be used to apply selective pulses to C2 [64].

Liquid state NMR involves bulk systems; typically, on the order of 10^{15} sample molecules occur in the sample being examined. The signal which is read out from the sample is an *ensemble average* over all those molecules, not a projective measurement which yields a single result, as in the quantum circuit model. In an NMR machine, magnetic pick-up coils are used to determine the magnetization in the x-y plane. The signal read-out from the coils is then Fourier transformed to give a spectrum for the system. The number of observables whose ensemble average can be directly observed in this way is thus rather limited. However, by making use of *reading pulses* immediately before the final measurement, it is possible to greatly extend the range of observables which can be determined. For example, a $\pi/2$ rotation about the y axis takes a Z operator to an X operator. Thus, although the ensemble average $\langle Z\rangle$ for a single nuclei cannot be directly observed, by applying a $\pi/2$ reading pulse about the y axis immediately before observation, the value of $\langle Z\rangle$ before the reading pulse can be inferred from the observed value of $\langle X\rangle$ after the reading pulse.

At room temperature, the initial state of the system is highly mixed. At temperature T, the spins start out in the state $\exp(-H/kT)$, where k is Boltzmann's constant, H is the Hamiltonian,

and a normalization factor, the partition function, has been omitted. To a first approximation, the coupling between nuclear spins may be omitted, and at high temperature the state of the system is proportional to

$$\prod_i \exp(-\omega_i Z_i/kT) \approx \prod_i (I - \omega_i Z_i/kT), \tag{2.37}$$

which is a mixture of computational basis states. This state does not appear to be at all like the pure computational basis state which is used in the quantum circuit model of quantum computation. There is a clever idea which allows us to work around this problem, suggested independently by Cory, Fahmy and Havel [51], and Gershenfeld and Chuang [69]. The idea is to extract a part of the state of the system which "looks like" a pure state. Perhaps the simplest scheme to illustrate the basic idea is the following method, known as *temporal labeling* [96].

Suppose we have a molecule with n nuclei. The idea is to define a set of unitary operators which permute all the computational basis states, $|0\rangle, \ldots, |2^n - 2\rangle$ amongst themselves, while leaving the state $|2^n - 1\rangle$ alone. We can then perform a series of $2^n - 1$ experiments as follows. In each experiment, the corresponding unitary operator is applied before the experiment begins. At the end, the experimental results from all $2^n - 1$ experiments are averaged. The net contribution due to the states $|0\rangle, \ldots, |2^n - 2\rangle$ averages out, leaving a net contribution due only to the state $|2^n - 1\rangle$. Thus we have performed a computation with an *effectively* pure state.

To see how this works in more detail, define unitary operators U_k, $0 \leq k \leq 2^n - 2$, by $U_k|x\rangle \equiv |x + k\rangle$ for $1 \leq x \leq 2^n - 2$ and $U_k|2^n - 1\rangle = |2^n - 1\rangle$, where the addition is done modulo $2^n - 1$. It is straightforward to efficiently implement such operations using standard quantum gates [7, 11], so this can be done in NMR. Note then that if $\rho = \sum_x p_x |x\rangle\langle x|$ is diagonal in the computational basis then

$$\begin{aligned} \sum_k U_k \rho U_k^\dagger &= (2^n - 1) p_N |N\rangle\langle N| + (1 - p_N) \sum_{x \neq N} |x\rangle\langle x| & (2.38) \\ &= (2^n p_N - 1)|N\rangle\langle N| + (1 - p_N) I, & (2.39) \end{aligned}$$

where $N \equiv 2^n - 1$. Suppose in each of these experiments we perform the unitary U_k, followed by some unitary operation U, and then observe some component of the spin, say $\langle X_i \rangle$. Summing over the results observed in each of the $2^n - 1$ experiments, we obtain

$$\sum_{k=0}^{N-1} \mathrm{tr}(X_i U U_k \rho U_k^\dagger U^\dagger) = (2^n p_N - 1)\mathrm{tr}(X_i U |N\rangle\langle N| U^\dagger), \tag{2.40}$$

as $\mathrm{tr}(X_i I) = 0$. That is, the summed averages behave as if the pure state $|N\rangle\langle N|$ had been prepared, the unitary operation U applied to that pure state, and the average of X_i observed. Similar remarks apply to other observations which may be made in NMR. By appropriate labeling we can ensure that p_N is the smallest (or largest) of the the eigenvalues of the initial density operator, in which case $2^n p_N \neq 1$, unless we are at infinite temperature, and the ensemble is uniform. Even the small deviation away from uniformity available at room temperature can be exploited to make the factor $2^n p_N - 1$ appearing in front of the observed average large enough that this method can be used to obtain effectively pure state behaviour out of a mixed state system.

This method is known as temporal averaging because it requires that the experiment be repeated many different times, and the results summed. Temporal averaging is only one possible means for performing state preparation in NMR quantum information processing. It is an especially easy method to explain, but in the laboratory other methods may be considerably better. In our

experiments, a technique based upon the use of *gradient pulses* was used. The precise details of what was done are beyond our present scope, but the basic idea may be explained quite easily.

Essentially what is done is to vary the strength of the magnetic field applied in the z direction across the sample. This causes nuclei at different locations in the sample to rotate around the z axis at different frequencies. When applied for the appropriate length of time, the ensemble averaged values for the X and Y components of magnetization average to zero. That is, a gradient pulse applied to a single spin has the effect of setting the x and y components of the Bloch vector for the ensemble to zero, while leaving the z component of the Bloch vector untouched. Cory *et al* [51] have described how a combination of gradient pulses, rf pulses, and delays may be combined to prepare effectively pure states, along similar lines to the temporal labeling method described above. We will not give further details of this method here.

NMR-based approaches to quantum information processing have many attractive features. NMR is a well-developed technology, and a considerable amount of high quality, easy-to-use equipment has been developed for use off-the-shelf. The noise timescale is typically on the order of a second, while the time to perform a two qubit gate is on the order of one to ten milliseconds, giving a best-case estimate of about one thousand couplings possible, although there is no doubt that achieving this in a useful computation will be extraordinarily difficult. Present experimental work in NMR quantum information processing usually involves on the order of ten couplings.

With regard to the power of NMR quantum information processing from the point of view of computational complexity, and in comparison with the quantum circuit model, I will not essay an opinion here. A considerable amount of interesting discussion has taken place on or closely related to this topic and I refer the reader to, for example, [39, 69, 96, 100, 153] for further discussion. What does seem certain is that NMR provides a powerful means for conducting interesting investigations into small-scale quantum information processing. A few qubits may not be much, but it represents the current best we can do with our quantum information processors.

2.6.2 Experimental demonstration of quantum teleportation using NMR

The ideas of NMR quantum information processing have recently been exploited to provide an experimental demonstration of quantum teleportation, in collaboration with Knill and Laflamme [137]. The essential idea of the scheme is to implement the quantum circuit for teleportation discussed in section 2.3, using the NMR-based techniques for quantum information processing discussed in the previous subsection.

Our implementation of teleportation is performed using liquid state nuclear magnetic resonance (NMR), applied to an ensemble of molecules of labeled trichloroethylene TCE, as discussed in the previous section. To perform teleportation we make use of the Hydrogen nucleus (H), and the two Carbon 13 nuclei (C1 and C2), teleporting the state of the second Carbon nucleus to the Hydrogen. Figure 2.3 shows a schematic illustration of the teleportation process we used, based upon the circuit described in [31], illustrated in figure 2.1. The circuit has three inputs, which we will refer to as the *data* (C2), *ancilla* (C1), and *target* (H) qubits. The goal of the circuit is to teleport the state of the data qubit so that it ends up on the target qubit.

State preparation is done in our experiment using the gradient-pulse techniques described by Cory *et al* [51], and phase cycling [64, 74]. The unitary operations performed during teleportation may be implemented in a straightforward manner in NMR, using non-selective rf pulses tuned to the Larmor frequencies of the nuclear spins, and delays allowing entanglement to form through the interaction of neighboring nuclei, as described in the previous section. Commented pulse sequences for our experiment may be obtained on the world wide web [136].

An innovation in our experiment was the method used to implement the Bell basis measure-

ment. In NMR, the measurement step allows us to measure the expectation values of σ_x and σ_y for each spin, averaged over the ensemble of molecules, rather than performing a projective measurement in some basis. For this reason, we must modify the projective measurement step in the standard description of teleportation, while still preserving the remarkable teleportation effect.

We use a procedure inspired by Brassard *et al* [31], who suggested a two-part procedure for performing the Bell basis measurement. Part one of the procedure is to rotate from the Bell basis into the computational basis, $|00\rangle, |01\rangle, |10\rangle, |11\rangle$. We implement this step in NMR by using the natural spin-spin coupling between the Carbon nuclei, and rf pulses. Part two of the procedure is to perform a projective measurement in the computational basis. As Brassard *et al* point out, the effect of this two part procedure is equivalent to performing the Bell basis measurement, and leaving the data and ancilla qubits in one of the four states, $|00\rangle, |01\rangle, |10\rangle, |11\rangle$, corresponding to the different measurement results.

We cannot directly implement the second step in NMR. Instead, we exploit the natural phase decoherence occurring on the Carbon nuclei to achieve the same effect. Recall that phase decoherence completely randomizes the phase information in these nuclei and thus will destroy coherence between the elements of the above basis. Its effect on the state of the Carbon nuclei is to diagonalize the state in the computational basis,

$$\begin{aligned} \rho \quad \longrightarrow \quad & |00\rangle\langle 00|\rho|00\rangle\langle 00| + |01\rangle\langle 01|\rho|01\rangle\langle 01| + |10\rangle\langle 10|\rho|10\rangle\langle 10| \\ & + |11\rangle\langle 11|\rho|11\rangle\langle 11|. \end{aligned} \tag{2.41}$$

As emphasized by Zurek [206], the decoherence process is indistinguishable from a measurement in the computational basis for the Carbons accomplished by the environment. We do not observe the result of this measurement explicitly, however the state of the nuclei selected by the decoherence process contains the measurement result, and therefore we can do the final transformation conditional on the particular state the environment has selected. As in the scheme of Brassard *et al*, the final state of the Carbon nuclei is one of the four states, $|00\rangle, |01\rangle, |10\rangle, |11\rangle$, corresponding to the four possible results of the measurement.

In our experiment, we exploit the natural decoherence properties of the TCE molecule. The phase decoherence times (T_2) for the C1 and C2 are approximately $0.4s$ and $0.3s$. All other T_2 and T_1 times for all three nuclei are much longer, with a T_2 time for the Hydrogen of approximately $3s$, and relaxation times (T_1) of approximately $20 - 30s$ for the Carbons, and $5s$ for the Hydrogen.

This implies that for delays on the order of $1s$, we can approximate the total evolution by exact phase decoherence on the Carbon nuclei. The total scheme therefore implements a measurement in the Bell basis, with the result of the measurement stored as classical data on the Carbon nuclei following the measurement. We can thus teleport the information from the Carbon to the Hydrogen and verify that the information in the final state decays at the Hydrogen rate and not the Carbon one.

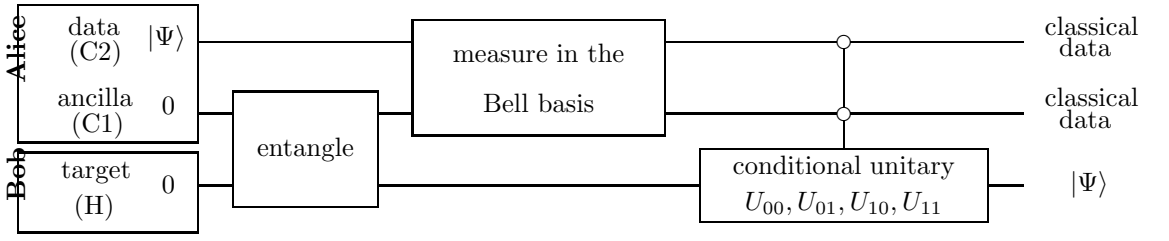

Figure 2.3: Schematic of quantum teleportation. See text for a full description.

Examining figure 2.3 we see how remarkable teleportation is from this point of view. During the stage labeled "Measure in the Bell basis" in figure 2.3, we allow the C1 and C2 nuclei to decohere and thus be measured by the environment, destroying all phase information on the data and ancilla qubits. Experimentally, the use of multiple refocusing pulses ensures that the data qubit has effectively not interacted with the target qubit. Classical intuition therefore tells us that the phase information about the input state, $|\Psi\rangle$, has been lost forever. Nevertheless, quantum mechanics predicts that we are still able to recover the complete system after this decoherence step, by quantum teleportation.

We implemented this scheme in TCE using a Bruker DRX-500 NMR spectrometer. Experimentally, we determined the Larmor and coupling frequencies for the Hydrogen, C1 and C2 to be:

$$\omega_H \approx 500.133491\text{MHz}; \qquad \omega_{C1} \approx 125.772580\text{MHz}; \qquad \omega_{C2} \approx \omega_{C1} - 911\text{Hz} \tag{2.42}$$

$$J_{H\,C1} \approx 201\text{Hz}; \qquad J_{C1\,C2} \approx 103\text{Hz}. \tag{2.43}$$

The coupling frequencies between H and C2, as well as the Chlorines to H, C1 and C2, are much lower, on the order of ten Hertz for the former, and less than a Hertz for the latter. Experimentally, these couplings are suppressed by multiple refocusings, and will be ignored in the sequel. Note that the frequencies of C1 and C2 are not identical; they have slightly different frequencies, due to the different chemical environments of the two atoms.

We performed two separate sets of experiments. In one set, the full teleportation process was executed, making use of a variety of decoherence delays in place of the measurement. The readout was performed on the Hydrogen nucleus, and a figure of merit – the dynamic fidelity – was calculated for the teleportation process. The dynamic fidelity is a quantity in the range 0 to 1 which measures the combined strength of *all* noise processes occurring during the process, which we will study in detail in Chapter 5[6]. In particular, an dynamic fidelity of 1 indicates perfect teleportation, while an dynamic fidelity of 0.25 indicates complete randomization of the state. Perfect *classical transmission* corresponds to an dynamic fidelity of 0.5, so dynamic fidelities of greater than 0.5 indicates that teleportation of some quantum information is taking place.

The second set of experiments was a control set. In those experiments, only the state preparation and initial entanglement of H and C1 were performed, followed by a delay for decoherence on C1 and C2. The readout was performed in this instance on C2, and once again, a figure of merit, the dynamic fidelity, was calculated for the entire process.

The results of our experiment are shown in figure 2.4, where the dynamic fidelity is plotted against the length of the delay which was used for the decoherence. Errors in our experiment arise from the strong coupling effect, imperfect calibration of rf pulses, and rf field inhomogeneities. The estimated uncertainty in our values for dynamic fidelity are less than ± 0.05. These uncertainties are due primarily to rf field inhomogeneity and imperfect calibration of rf pulses.

In order to determine the dynamic fidelities for the teleportation and control experiments, we performed *quantum process tomography*. This procedure, described in detail in section 3.4, exploits the linearity of quantum mechanics to completely characterize a quantum mechanical process. In particular, we will show in section 3.3 that the linearity of quantum mechanics implies that the single qubit input and output for the teleportation process are always related by a linear quantum operation. By preparing a complete set of four linearly independent initial states, we were able to obtain a complete description of the quantum process. In particular, we used a procedure known as *quantum state tomography* [181, 169] to determine the output states, and used the linearity of

[6] In the language of that Chapter, we determined the dynamic fidelity for teleportation of the state I/2.

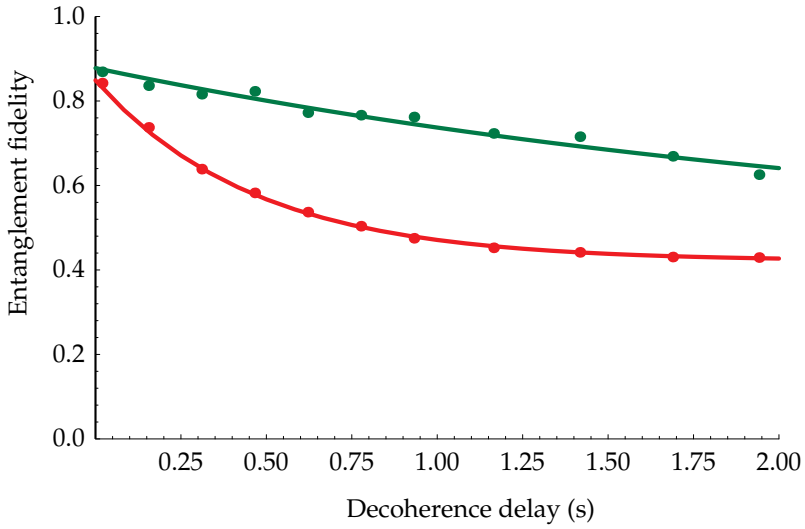

Figure 2.4: Dynamic fidelity is plotted as a function of decoherence time. The bottom curve is a control run where the information remains in C2. The curve shows a decay time of approximately 0.5s. The top curve represents the fidelity of the quantum teleportation process. The decay time is approximately 2.6s. The information is preserved for a longer time, corresponding approximately to the combined effects of decoherence and relaxation for the Hydrogen, confirming the prediction of teleportation.

quantum mechanics to extend this to a description of the complete process. This description, in turn, enabled us to calculate the dynamic fidelity for the process, as described in sections 3.4 and 5.3.

Three elements ought to be noted in figure 2.4. First, for small decoherence delays, the dynamic fidelity for the teleportation experiments significantly exceeds the value of 0.5 for perfect classical transmission of data, indicating that we have successfully teleported quantum information from C2 to H, with reasonable fidelity. Second, it is notable that the dynamic fidelity decays very quickly for the control experiments as the delay is increased. Theoretically, we expect this to be the case, due to a $T2$ time for C2 of approximately $0.3s$. Third, the decay of the dynamic fidelity for the teleportation experiments occurs much more slowly. Theoretically, we expect this decay to be due mainly to the effect of phase decoherence and relaxation on the *Hydrogen*. Our experimental observations are consistent with this prediction, and provide more support for the claim that the quantum data is being teleported in this set of experiments.

In conclusion, we have exhibited evidence of successful quantum teleportation in liquid state NMR. This experiment is not the first experimental implementation of quantum teleportation, however it is the first implementation in NMR. Earlier experiments by Boschi *et al* [28] and Bouwmeester *et al* [29] used optical methods to achieve quantum teleportation. The present NMR-based method illustrates some of the advantages of using NMR to do elementary quantum information processing. In particular, the NMR experiment was relatively straightforward to set up and perform, using off-the-shelf methods, and could be repeated easily in many laboratories around the world. By contrast, the optical techniques required much more customized equipment, and were generally much more

difficult to achieve.

Summary of Chapter 2: Quantum information: fundamentals

- **Qubits:** The fundamental unit of quantum information. A two level quantum system.
- **Superdense Coding:** Preshared entanglement can be used to transmit two classical bits with the transmission of only one qubit.
- **Quantum teleportation:** Preshared entanglement can be used to transmit a qubit with the transmission of two classical bits.
- **Quantum circuits:** *ad hoc* model with the following features:
  1. **Classical external control.**
  2. **A suitable state space:** n qubits.
  3. **Ability to prepare states in the computation basis.**
  4. **Dynamics:** An *algorithm* for applying *quantum gates* (controlled-not and single qubit unitary gates) and *projective measurements* in the *computational basis* to the system.
- **Experimental implementation of quantum teleportation in NMR**

Chapter 3

Quantum operations

Quantum mechanics describes the *dynamics* which can occur in physical systems. Elementary quantum mechanics texts usually do this by separating the dynamics into two different types. The first type is the evolution of a closed quantum mechanical system, which is assumed to be described by Schrödinger's equation. Under such an evolution, the change in the state $|\psi\rangle$ of a quantum system between two fixed times is described by a unitary operator U which depends on those times,

$$|\psi\rangle \rightarrow U|\psi\rangle. \tag{3.1}$$

The equivalent map on density operators is given by

$$\rho \rightarrow \mathcal{E}(\rho) \equiv U\rho U^\dagger. \tag{3.2}$$

The second type of dynamics described in basic quantum mechanics texts is associated with the *measurement* of a quantum mechanical system. The system being measured is no longer a closed system, since it is interacting with the measuring device. The usual way to describe such a measurement is the following. Suppose a measurement is performed which has outcomes labeled by m. Then to each outcome m there is associated a projector P_m onto the state space of the system such that

$$P_m P_n = \delta_{mn} P_m \tag{3.3}$$

$$\sum_m P_m = I. \tag{3.4}$$

If the state of the system immediately before the measurement was ρ, and the result of the measurement is m, then the state of the system immediately after the measurement is

$$\frac{\mathcal{E}_m(\rho)}{\mathrm{tr}(\mathcal{E}_m(\rho))}, \tag{3.5}$$

where

$$\mathcal{E}_m(\rho) \equiv P_m \rho P_m. \tag{3.6}$$

Moreover, the probability of obtaining this measurement result is given by

$$p(m) = \mathrm{tr}(\mathcal{E}_m(\rho)). \tag{3.7}$$

The maps $\mathcal{E}$ and $\mathcal{E}_m$ are examples of *quantum operations*. The theory of quantum operations can be used to describe a wide class of state changes that may occur in quantum systems. You may wonder how it is possible to go beyond the usual textbook description of state changes in terms of unitary transformations and projective measurements. The key observation is that many state changes of interest occur in *open quantum systems*. The interaction of the quantum system with an external world allows dynamics that are neither unitary nor described by the usual model of projective measurements.

To make the idea of quantum operations more concrete, consider the following example. Suppose we have a single qubit quantum system, the *principal* system, in a state ρ, which is brought into contact with an environment. We will suppose this environment is also a single qubit system, which is initially in the state $|0\rangle$. For instance, the principal system might be a nuclear spin in a molecule being used to do NMR, and the environment a neighbouring spin. Left to themselves these systems will interact according to some unitary interaction U. For the sake of definiteness we will suppose that U is the controlled not operation, with the principal system the control qubit, and the environment the data qubit. U can be written

$$U = P_0 \otimes I + P_1 \otimes X, \tag{3.8}$$

where the first system is the principal system, the second system is the environment, and $P_0 \equiv |0\rangle\langle 0|, P_1 \equiv |1\rangle\langle 1|$ are projectors. The state of the joint system after the interaction is

$$\begin{aligned} &U(\rho \otimes |0\rangle\langle 0|)U^\dagger = \\ &P_0\rho P_0 \otimes |0\rangle\langle 0| + P_1\rho P_1 \otimes |1\rangle\langle 1| + P_0\rho P_1 \otimes |0\rangle\langle 1| + P_1\rho P_0 \otimes |1\rangle\langle 0|. \end{aligned} \tag{3.9}$$

Tracing out the environment we obtain the final state of the principal system,

$$P_0\rho P_0 + P_1\rho P_1. \tag{3.10}$$

That is, the evolution of the principal system can be described by the map $\mathcal{E}$,

$$\rho \rightarrow \mathcal{E}(\rho) \equiv P_0\rho P_0 + P_1\rho P_1. \tag{3.11}$$

The map just described is an example of a quantum operation, in which the quantum state undergoes one single, definite evolution. By contrast, in the case of a measurement, several different outcomes may occur, each outcome being associated with a particular classical measurement value. For example, suppose a principal system consisting of one qubit is being coupled, once again, to a one qubit environment. The initial state of the total system is again assumed to be $\rho \otimes |0\rangle\langle 0|$, and the coupling is described by a unitary operator U defined by

$$U = \frac{X}{\sqrt{2}} \otimes I + \frac{Y}{\sqrt{2}} \otimes X. \tag{3.12}$$

Following the unitary evolution, a measurement is done on the environment qubit, in the computational basis. Note that the state of the system after the unitary evolution is

$$\frac{X\rho X \otimes |0\rangle\langle 0| + Y\rho Y \otimes |1\rangle\langle 1| + X\rho Y \otimes |0\rangle\langle 1| + Y\rho X \otimes |1\rangle\langle 0|}{2}. \tag{3.13}$$

Conditioned on the result of the measurement, we see by inspection that the state of the principal system after the measurement is either $X\rho X$ or $Y\rho Y$, depending upon whether the measurement result was 0 or 1, with probability 1/2 for each of the two possibilities. Again, these are quantum operations, this time associated with different measurement outcomes possible in the process.

The primary purpose of this Chapter is to review the general theory of quantum operations. In addition to elementary review material, the Chapter shows how the quantum operations formalism can be used to gain insight into quantum teleportation, and describes *quantum process tomography*, a general method for the experimental determination of the dynamics of a quantum system. The elementary material appearing here has its origins in earlier work by people such as Hellwig and Kraus [79, 80], Choi [36] and Kraus [102]. The material relating quantum teleportation and the quantum operations formalism is based upon a collaboration with Caves [135], and the work on quantum process tomography is based upon a collaboration with Chuang [40]. In places the Chapter contains rather detailed mathematics; upon a first read, these sections may be read lightly, and returned to later for reference purposes.

3.1 Quantum operations: fundamentals

Suppose we have a quantum system Q, initially in an *input state*, ρ. We suppose some physical process occurs, which results in an *output state*, ρ'. That output state need not even be a state of the same system; all we require is that the final state is uniquely determined by some physical process, starting with the input state. What requirements must the map $\mathcal{E} : \rho \rightarrow \rho'$ satisfy? We will enumerate a set of axioms which any such map must satisfy, and then go on to show that any map satisfying these requirements is physically reasonable.

The formalism we develop shall, ideally, include deterministic quantum processes, such as unitary evolution and interaction with an inaccessible environment, as well as measurements, in which a quantum system undergoes a state change chosen at random, depending on what measurement outcome occurred.

To cope with the case of measurements, it turns out that it is extremely convenient to make the convention that the map $\mathcal{E} : \rho \rightarrow \rho'$ does not necessarily preserve the trace property of density operators, that $\mathrm{tr}(\rho) = 1$. Rather, we make the convention that $\mathcal{E}$ is to be defined in such a way that $\mathrm{tr}(\mathcal{E}(\rho))$ is equal to the probability of the measurement outcome described by $\mathcal{E}$ occurring. More concretely, suppose that we are doing a projective measurement in the computational basis of a single qubit. Then two quantum operations are used to describe this process, defined by $\mathcal{E}_0(\rho) \equiv |0\rangle\langle 0|\rho|0\rangle\langle 0|$ and $\mathcal{E}_1(\rho) \equiv |1\rangle\langle 1|\rho|1\rangle\langle 1|$. Notice that the probabilities of the respective outcomes are correctly given by $\mathrm{tr}(\mathcal{E}_0(\rho))$ and $\mathrm{tr}(\mathcal{E}_1(\rho))$. With this convention the correctly normalized final quantum state is therefore

$$\frac{\mathcal{E}(\rho)}{\mathrm{tr}(\mathcal{E}(\rho))}. \tag{3.14}$$

Thus, generically, we impose a requirement of mathematical convenience, that $\mathrm{tr}(\mathcal{E}(\rho))$ be equal to the probability of the process represented by $\mathcal{E}$ occurring, when ρ is the initial state. In the case where the process is deterministic, that is, no measurement is taking place, this reduces to the requirement that $\mathrm{tr}(\mathcal{E}(\rho)) = 1 = \mathrm{tr}(\rho)$. In this case, we say that the quantum operation is a *complete* quantum operation, since on its own it provides a complete description of the quantum process. On the other hand, if there is a ρ such that $\mathrm{tr}(\mathcal{E}(\rho)) < 1$, then the quantum operation is *incomplete*, since on its own it does not provide a complete description of the processes that may occur in the system. A *physical* quantum operation is one that satisfies the requirement that probabilities never exceed 1, $\mathrm{tr}(\mathcal{E}(\rho)) \leq 1$.

Next, we impose our first physical requirement on quantum operations. Suppose the input ρ to the quantum operation is obtained by randomly selecting the state from an ensemble $\{p_i, \rho_i\}$ of quantum states, that is, $\rho = \sum_i p_i \rho_i$. Then we would expect that the resulting state, $\mathcal{E}(\rho)/\mathrm{tr}(\mathcal{E}(\rho))$

corresponds to a random selection from the ensemble $\{p(i|\mathcal{E}), \mathcal{E}(\rho_i)/\mathrm{tr}(\mathcal{E}(\rho_i))\}$, where $p(i|\mathcal{E})$ is the probability that the state prepared was ρ_i, given that the process represented by $\mathcal{E}$ occurred. Thus, we demand that

$$\mathcal{E}(\rho) = p(\mathcal{E}) \sum_i p(i|\mathcal{E}) \frac{\mathcal{E}(\rho_i)}{\mathrm{tr}(\mathcal{E}(\rho_i))}, \tag{3.15}$$

where $p(\mathcal{E}) = \mathrm{tr}(\mathcal{E}(\rho))$ is the probability that $\mathcal{E}$ occurs on input of ρ. By Bayes' rule,

$$p(i|\mathcal{E}) = p(\mathcal{E}|i) \frac{p_i}{p(\mathcal{E})} = \frac{\mathrm{tr}(\mathcal{E}(\rho_i))p_i}{p(\mathcal{E})}, \tag{3.16}$$

so the equation (3.15) reduces to

$$\mathcal{E}(\sum_i p_i \rho_i) = \sum_i p_i \mathcal{E}(\rho_i). \tag{3.17}$$

That is, we require that quantum operations be *convex linear* on the set of density operators. Indeed, any convex linear map on density operators can be uniquely extended to a linear map on Hermitian operators, so we make this additional requirement, again, as a mathematical convenience.

Finally, we require that the quantum operation must preserve the positivity of density operators. This requirement, known as *complete positivity*, means that quantum operations take positive operators to positive operators. This requirement applies both to density operators on the system for which the dynamics is occurring, the *principal system*, and also for super-systems of the principal system.

To illustrate the importance of this point, consider the transpose operation on a single qubit. By definition, this map transposes the density operator in the computational basis:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \xrightarrow{T} \begin{bmatrix} a & c \\ b & d \end{bmatrix}. \tag{3.18}$$

This map preserves positivity of a single qubit. However, suppose that qubit is part of a two qubit system initially in the entangled state

$$\frac{|00\rangle + |11\rangle}{\sqrt{2}}, \tag{3.19}$$

and the transpose operation is applied to the first of these two qubits, while the second qubit is subject to trivial dynamics. Then the density operator of the system after the dynamics has been applied is

$$\frac{1}{2} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{3.20}$$

It is easy to verify that this operator has eigenvalues $1/2, 1/2, 1/2$ and $-1/2$, so this is not a valid density operator. Thus, the transpose operator is an example of a map that preserves the positivity of operators on the principal system, but does not continue to preserve positivity when applied to systems which contain the principal system as a subsystem[1].

Summarizing, the requirements for a map to be a quantum operation are as follows:

[1] According to Weinberg [184] there are selection rules in some system that prohibit, for example, superpositions of states with different electric charge existing. It is amusing to speculate that in systems in which such selection rules exist it might be allowable for systems to undergo dynamics which are not completely positive, as this would not necessarily lead to density operators which were not positive, and thus unphysical.

1. By definition, $\mathrm{tr}(\mathcal{E}(\rho))$ is the probability that the process represented by $\mathcal{E}$ occurs, when ρ is the initial state.

2. The map $\mathcal{E}$ is a linear map. The domain of $\mathcal{E}$ is the (real vector space) of Hermitian operators on H_Q, the input Hilbert space. The range of $\mathcal{E}$ is contained in the (real vector space) of Hermitian operators on H'_Q, the output Hilbert space.

3. The map $\mathcal{E}$ is *completely positive*. That is, suppose we introduce a second system, R, with (finite dimensional) Hilbert space H_R. Let $\mathcal{I}$ denote the identity map on system R. Then the map $\mathcal{I} \otimes \mathcal{E}$ takes positive operators to positive operators.

Surprisingly to me, at least, these requirements are sufficient to characterize quantum operations. Later, we will show how any map satisfying these requirements can be physically realized, in a finite dimensional quantum system. One step along the way to this result is an elegant representation theorem which relates these abstract requirements for a quantum operation to an explicit formula:

Theorem 1 ***(Operator-sum representation)*** *[102]*
The map $\mathcal{E}$ is a quantum operation if and only if

$$\mathcal{E}(A) = \sum_i E_i A E_i^\dagger, \tag{3.21}$$

for some set of operators E_i which map the input Hilbert space to the output Hilbert space.

The operators E_i appearing in this expression are said to generate an *operator-sum* representation for the quantum operation $\mathcal{E}$.

Proof [157]

Suppose $\mathcal{E}(A) = \sum_i E_i A E_i^\dagger$. $\mathcal{E}$ is obviously linear, so to check that $\mathcal{E}$ is a quantum operation we need only prove that it is completely positive. Let A be any positive operator acting on the state space of an extended system, RQ, and let $|\psi\rangle$ be some state of that system. Defining $|\phi_i\rangle \equiv (I_R \otimes E_i^\dagger)|\psi\rangle$, we have

$$\begin{aligned} \langle\psi|(I_R \otimes E_i)A(I_R \otimes E_i^\dagger)|\psi\rangle &= \langle\phi_i|A|\phi_i\rangle & (3.22)\\ &\geq 0, & (3.23)\end{aligned}$$

by the positivity of the operator A. It follows that

$$\langle\psi|(\mathcal{I} \otimes \mathcal{E})(A)|\psi\rangle = \sum_i \langle\phi_i|A|\phi_i\rangle \geq 0, \tag{3.24}$$

and thus for any positive operator A, the operator $(\mathcal{I} \otimes \mathcal{E})(A)$ is also positive, as required. This completes the first part of the proof.

Suppose next that $\mathcal{E}$ is a quantum operation. Our aim will be to find an operator-sum representation for $\mathcal{E}$. Suppose we introduce a system, R, with the same dimension as the original quantum system, Q. Let $|i_R\rangle$ and $|i_Q\rangle$ be orthonormal bases for R and Q. It will be convenient to use the same index, i, for these two bases, and this can certainly be done as R and Q have the same dimensionality. Define a joint state $|\alpha\rangle$ of RQ by

$$|\alpha\rangle \equiv \sum_i |i_R\rangle|i_Q\rangle. \tag{3.25}$$

The state $|\alpha\rangle$ is, up to a normalization factor, a maximally entangled state of the systems R and Q. This interpretation of $|\alpha\rangle$ as a maximally entangled state may help in understanding the following construction. Next, we define an operator σ on the state space of RQ by

$$\sigma \equiv (\mathcal{I}_R \otimes \mathcal{E})(|\alpha\rangle\langle\alpha|). \tag{3.26}$$

We may think of this as the result of applying the quantum operation $\mathcal{E}$ to one half of a maximally entangled state of the system RQ. It is a truly remarkable fact, which we will now demonstrate, that the operator σ completely specifies the quantum operation $\mathcal{E}$. That is, to know how $\mathcal{E}$ acts on an arbitrary state of Q, it is sufficient to know how it acts on a single maximally entangled state of Q with another system[2].

The trick which allows us to recover $\mathcal{E}$ from σ is as follows. Let $|\psi\rangle = \sum_j \psi_j |j_Q\rangle$ be any state of system Q. Define a corresponding state $|\tilde{\psi}\rangle$ of system R by the equation

$$|\tilde{\psi}\rangle \equiv \sum_j \psi_j^* |j_R\rangle. \tag{3.27}$$

Notice that

$$\begin{aligned}
\langle\tilde{\psi}|\sigma|\tilde{\psi}\rangle &= \langle\tilde{\psi}| \left(\sum_{ij} |i_R\rangle\langle j_R| \otimes \mathcal{E}(|i_Q\rangle\langle j_Q|) \right) |\tilde{\psi}\rangle && (3.28)\\
&= \sum_{ij} \psi_i \psi_j^* \mathcal{E}(|i_Q\rangle\langle j_Q|) && (3.29)\\
&= \mathcal{E}(|\psi\rangle\langle\psi|). && (3.30)
\end{aligned}$$

Let $\sigma = \sum_i |s_i\rangle\langle s_i|$ be some decomposition of σ, where the vectors $|s_i\rangle$ need not be normalized. Define a map

$$E_i(|\psi\rangle) \equiv \langle\tilde{\psi}|s_i\rangle. \tag{3.31}$$

A little thought shows that this map is a linear map, so E_i is a linear operator on the state space of Q. Furthermore, we have

$$\begin{aligned}
\sum_i E_i|\psi\rangle\langle\psi|E_i^\dagger &= \langle\tilde{\psi}|s_i\rangle\langle s_i|\tilde{\psi}\rangle && (3.32)\\
&= \langle\tilde{\psi}|\sigma|\tilde{\psi}\rangle && (3.33)\\
&= \mathcal{E}(|\psi\rangle\langle\psi|). && (3.34)
\end{aligned}$$

Thus

$$\mathcal{E}(|\psi\rangle\langle\psi|) = \sum_i E_i|\psi\rangle\langle\psi|E_i\dagger, \tag{3.35}$$

for all pure states, $|\psi\rangle$, of Q. By linearity it follows that

$$\mathcal{E}(A) = \sum_i E_i A E_i^\dagger \tag{3.36}$$

[2]It is interesting and enlightening to contemplate a similar construction for classical systems, based upon a maximally *correlated* state of two classical systems. A construction of this sort is given at the beginning of Chapter 5.

in general.

QED

This result allows us to give easy proofs that many interesting maps are quantum operations. For instance, it is clear that the unitary evolution $\mathcal{E}(\rho) = U\rho U^\dagger$ is a quantum operation. It is also clear that a measurement is described by a set of quantum operations $\mathcal{E}_m(\rho) = P_m \rho P_m$ indexed by the measurement outcome m.

Slightly less obviously, we see that the trace map $A \rightarrow \mathrm{tr}(A)$ is a quantum operation. To see this, let H_Q be any input Hilbert space, spanned by an orthonormal basis $|1\rangle \ldots |d\rangle$, and let H'_Q be a one dimensional output space, spanned by the state $|0\rangle$. Define

$$\mathcal{E}(A) \equiv \sum_{i=1}^{d} |0\rangle\langle i|A|i\rangle\langle 0|, \tag{3.37}$$

so that $\mathcal{E}$ is a quantum operation, by the operator-sum representation theorem. Notice that $\mathcal{E}(A) = \mathrm{tr}(A)|0\rangle\langle 0|$, so that, up to the unimportant $|0\rangle\langle 0|$ multiplier, this quantum operation is identical to the trace function.

An even more useful result is the observation that the partial trace is a quantum operation. Suppose we have a joint system AB, and wish to trace out system B. Let $|j\rangle$ be a basis for system B. Define a linear operator $E_i : H_{AB} \rightarrow H_A$ by

$$E_i \sum_j \lambda_j |a_j\rangle|j\rangle \equiv \lambda_i |a_i\rangle, \tag{3.38}$$

where λ_j are complex numbers, and $|a_j\rangle$ are arbitrary states of system A. Define

$$\mathcal{E}(A) \equiv \sum_i E_i A E_i^\dagger. \tag{3.39}$$

By the operator-sum representation theorem for quantum operations, this is a quantum operation from system AB to system A. Notice that

$$\mathcal{E}(A \otimes |j\rangle\langle j'|) = A\delta_{j,j'} = \mathrm{tr}_B(A \otimes |j\rangle\langle j'|), \tag{3.40}$$

where A is any Hermitian operator on the state space of system A, and $|j\rangle$ and $|j'\rangle$ are members of the orthonormal basis for system B. By linearity of $\mathcal{E}$ and tr_B, it follows that $\mathcal{E} = \mathrm{tr}_B$.

In terms of the operator-sum representation, it is easy to characterize a quantum operation as being complete, incomplete, or physical. Recall that a quantum operation is complete if $\mathrm{tr}(\mathcal{E}(\rho)) = 1$ for all input states ρ. Clearly, this is equivalent to the requirement that $\sum_i E_i^\dagger E_i = I$ for the operators E_i in the operator-sum representation. Similarly, the property that a quantum operation be incomplete is equivalent to the condition that $\sum_i E_i^\dagger E_i < I$, while the property that a quantum operation is physical is equivalent to the condition that $\sum_i E_i^\dagger E_i \leq I$.

One reason for our interest in the operator-sum representation is that it gives us a way of characterizing the dynamics of a system in terms of *intrinsic quantities*. Non-unitary behaviour of quantum system can only arise because of the action of external systems. The operator sum representation gives us a way of describing the dynamics of the principal system, without having to explicitly consider properties of those external systems; all that we need to know is bundled up into the operators E_i, which act on the Hilbert space of the principal system alone. Furthermore, we will see soon that many different interactions with an external system may give rise to the same dynamics on the principal system. If it is only the dynamics of the principal system which are

of interest then it makes sense to choose a representation of the dynamics which does not include unimportant information about other systems.

We can relate the operator-sum representation picture of quantum operations to the idea of a quantum system interacting with other systems. We will prove two results. The first result shows how to determine the operator-sum representation appropriate for a quantum system interacting in a specified way with other quantum systems. The second result shows that for *any* quantum operation, we can always find a reasonable *model external system and dynamics* which give rise to that quantum operation. By reasonable, we here mean that the dynamics must be either a unitary evolution or a projective measurement.

Suppose we have a quantum system initially in a state ρ. We will denote this system by the letter Q. Adjoined to Q is another system which we will refer to variously as the *ancilla* or *environment* system, and denote by E. We suppose that Q and E are initially independent systems, and that E starts in some standard state, σ. The joint state of the system is thus initially

$$\rho^{QE} = \rho \otimes \sigma. \tag{3.41}$$

We suppose that the systems interact according to some unitary interaction U.

After the unitary interaction a measurement may be performed on the joint system. This measurement is described by projectors P_m. The case where no measurement is made corresponds to the special case where there is only a single measurement outcome, $m = 0$, which corresponds to the projector $P_0 \equiv I$.

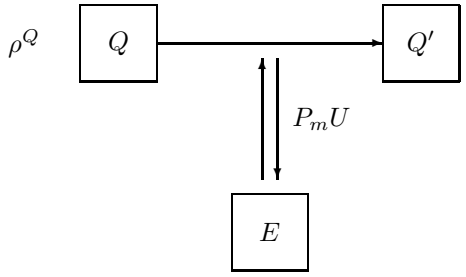

Figure 3.1: Environmental model for a quantum operation.

The situation is summarized in figure 3.1. Our aim is to determine the final state of Q as a function of the initial state, ρ. The final state of QE is given by

$$\frac{P_m U(\rho \otimes \sigma) U^\dagger P_m}{\mathrm{tr}(P_m U(\rho \otimes \sigma) U^\dagger P_m)}, \tag{3.42}$$

given that measurement outcome m occurred. Tracing out E we see that the final state of Q alone is

$$\frac{\mathrm{tr}_E(P_m U(\rho \otimes \sigma) U^\dagger P_m)}{\mathrm{tr}(P_m U(\rho \otimes \sigma) U^\dagger P_m)}. \tag{3.43}$$

This representation of the final state involves the initial state σ of the environment, the interaction U and the measurement operators P_m. Define a map

$$\mathcal{E}_m(\rho) \equiv \mathrm{tr}_E(P_m U(\rho \otimes \sigma) U^\dagger P_m). \tag{3.44}$$

Note that $\mathrm{tr}(\mathcal{E}_m(\rho))$ is the probability of outcome m of the measurement occurring. Let $\sigma = \sum_j q_j |j\rangle\langle j|$ be an ensemble decomposition for σ. Introduce an orthonormal basis $|k\rangle\langle k|$ for the system E. Note that

$$\begin{aligned} \mathrm{tr}(\mathcal{E}_m(\rho)) &= \sum_{jk} q_j \mathrm{tr}_E(|k\rangle\langle k|P_m U(\rho \otimes |j\rangle\langle j|)U^\dagger P_m|k\rangle\langle k|) & (3.45) \\ &= \sum_{jk} E_{jk} \rho E_{jk}^\dagger, & (3.46) \end{aligned}$$

where

$$E_{jk} \equiv \sqrt{q_j}\langle k|P_m U|j\rangle. \tag{3.47}$$

This equation gives an explicit means for calculating the operators appearing in an operator-sum representation for $\mathcal{E}_m$, given that the initial state σ of E is known, and the dynamics between Q and E are known. Indeed, two examples of this prescription in action were already given, in the opening section to this Chapter.

We now review a construction converse to this, which shows that for any quantum operation $\mathcal{E}$, we can *mock up* the dynamics $\mathcal{E}$ using an appropriate model. The construction will only be given for quantum operations mapping the input space to the same output space, although it is mainly a matter of notation to generalize the construction to the more general case. In particular, we show that for any physical quantum operation, $\mathcal{E}$, there exists a model environment, E, starting in a pure state $|0\rangle$, and model dynamics specified by a unitary operator U and projector P onto E such that

$$\mathcal{E}(\rho) = \mathrm{tr}_E(PU(\rho \otimes |0\rangle\langle 0|)U^\dagger P). \tag{3.48}$$

To see this, suppose first that $\mathcal{E}$ is a complete quantum operation, with operator-sum representation generated by operators E_i satisfying the completeness relation $\sum_i E_i^\dagger E_i = I$, so we are only attempting to find an appropriate unitary operator U to model the dynamics. Let $|i\rangle$ be an orthonormal basis set for E, in one-to-one correspondence with the index i for the operators E_i. Define an operator U which has the following action on states of the form $|\psi\rangle|0\rangle$,

$$U|\psi\rangle|0\rangle \equiv \sum_i E_i|\psi\rangle|i\rangle. \tag{3.49}$$

Note that for arbitrary states $|\psi\rangle$ and $|\phi\rangle$ of Q,

$$\langle\psi|\langle 0|U^\dagger U|\phi\rangle|0\rangle = \sum_i \langle\psi|E_i^\dagger E_i|\phi\rangle = \langle\psi|\phi\rangle, \tag{3.50}$$

by the completeness relation. Thus the operator U can be extended to a unitary operator acting on the entire state space of the joint system. It is easy to verify that

$$\mathrm{tr}_E(U(\rho \otimes |0\rangle\langle 0|)U^\dagger) = \sum_i E_i \rho E_i^\dagger, \tag{3.51}$$

so this model provides a realization of the quantum operation $\mathcal{E}$.

Incomplete quantum operations can easily be modeled using a construction along the same lines. Simply introduce an extra operator, E_∞, into the set of operators E_i, chosen so that when summing over the complete set of i, including $i = \infty$, one obtains $\sum_i E_i^\dagger E_i = I$. Now repeat the same construction as before to obtain a unitary operator U. Following the unitary U, however, it is necessary to do a projection onto the states $|i\rangle$ where $i \neq \infty$, to remove this operator from the operator-sum representation of the quantum operation being modeled.

A more interesting generalization of this construction is the case of a set of physical quantum operations, $\{\mathcal{E}_m\}$. Note that if a set of quantum operations $\{\mathcal{E}_m\}$ corresponded to possible outcomes from a measurement, then the quantum operation $\sum_m \mathcal{E}_m$ is complete, since the probabilities of the distinct outcomes sum to one, $1 = \sum_m p(m) = (\sum_m \mathcal{E}_m)(\rho)$ for all possible inputs ρ.

Conversely, if we are given a set of physical quantum operations $\{\mathcal{E}_m\}$ such that $\sum_m \mathcal{E}_m$ then it is possible to construct a *measurement model* giving rise to this set of quantum operations. For each m, let E_{mi} be a set of operators generating an operator-sum representation for $\mathcal{E}_m$. Introduce an environmental system, E, with an orthonormal basis $|m, i\rangle$ in one-to-one correspondence with the set of indices for the operators generating the respective operator-sum representations. Analogously to the earlier construction, define an operator U such that

$$U|\psi\rangle|0\rangle = \sum_{mi} E_{mi}|\psi\rangle|m, i\rangle. \tag{3.52}$$

As before, this operator may be extended to a unitary operation, because of the completeness relation $\sum_{mi} E_{mi}^\dagger E_{mi} = I$. Next, define projectors $P_m \equiv \sum_i |m, i\rangle\langle m, i|$ on the environmental system, E.

Suppose we perform the unitary operation U on the state $\rho \otimes |0\rangle\langle 0|$, and follow that up with a measurement on the environmental system, with the measurement being defined by the complete set of orthogonal projectors P_m. Then it is easy to verify that the (unnormalized) state of the principal system if the measurement result m is recorded is $\sum_i E_{mi} \rho E_{mi}^\dagger = \mathcal{E}_m(\rho)$, with the probability of the outcome m being given by the trace of $\mathcal{E}_m(\rho)$, exactly as required.

3.1.1 Quantum operations on a single qubit

There is a nice geometric way of picturing quantum operations when the principal system is a single qubit. This method allows one to get an intuitive feel for the behaviour of quantum operations in terms of their action on the Bloch sphere. Recall from section 2.1 that the state of a single qubit can always be written in the Bloch representation,

$$\rho = \frac{I + \vec{\lambda} \cdot \vec{\sigma}}{2}, \tag{3.53}$$

where $\vec{\lambda}$ is a three component real vector.

In this representation, it turns out that an arbitrary complete quantum operation is equivalent to a map of the form

$$\vec{\lambda} \xrightarrow{\mathcal{E}} \vec{\lambda}' = M\vec{\lambda} + \vec{c}, \tag{3.54}$$

where M is a 3×3 matrix, and $\vec{c}$ is a constant vector. This is an *affine map*, mapping the Bloch sphere into itself. Suppose the operators E_i generating the operator-sum representation for $\mathcal{E}$ are written in the form

$$E_i = \alpha_i I + \sum_{k=1}^{3} a_{ik}\sigma_k. \tag{3.55}$$

Then it is not difficult to check that

$$\begin{aligned} M_{jk} &= \sum_l \left[a_{lj} a_{lk}^* + a_{lj}^* a_{lk} + \left(|\alpha_l|^2 - \sum_p a_{lp} a_{lp}^* \right) \delta_{jk} \right. \\ &\quad \left. + i \sum_p \epsilon_{jkp} (\alpha_l a_{lp}^* - \alpha_l^* a_{lp}) \right] \qquad (3.56) \\ c_k &= 2i \sum_l \sum_{jp} \epsilon_{jpk} a_{lj} a_{lp}^*, \qquad (3.57) \end{aligned}$$

where we have made use of the completeness relation $\sum_i E_i^\dagger E_i = I$ to simplify the expression for $\vec{c}$.

The meaning of the affine map equation (3.54) is made clearer by considering the polar decomposition [86] of the matrix M. Any real matrix M can always be written in the form

$$M = OS, \qquad (3.58)$$

where O is a real orthogonal matrix with determinant 1, representing a proper rotation, and S is a real symmetric matrix. Viewed this way, equation (3.54) is just a deformation of the Bloch sphere along principal axes determined by S, followed by a proper rotation due to O, followed by a displacement due to $\vec{c}$.

This picture can be used to obtain simple pictures of quantum operations on single qubits. For example, unitary operations correspond to (possibly improper) rotations of the Bloch sphere.

Less trivially, consider the completely decohering quantum operation,

$$\rho \rightarrow \mathcal{E}(\rho) = P_0 \rho P_0 + P_1 \rho P_1, \qquad (3.59)$$

for which we introduced an environmental model in the opening section of this Chapter. Using the above prescription it is easy to see that the corresponding map on the Bloch sphere is given by

$$(\lambda_x, \lambda_y, \lambda_z) \rightarrow (0, 0, \lambda_z). \qquad (3.60)$$

Geometrically, the Bloch vector is projected along the z axis, and the x and y components of the the Bloch vector are lost. This geometric picture makes it very easy to verify certain facts about this quantum operation. For example, it is easy to verify that the quantity $\mathrm{tr}(\rho^2)$ for a single qubit is equal to $(1+|\lambda|^2)/2$, where $|\lambda|$ is the norm of the Bloch vector. The projection process above cannot increase the norm of the Bloch vector, and therefore we can immediately conclude that $\mathrm{tr}(\rho^2)$ can only ever decrease for the completely decohering quantum operation. This is but one example of the use of this geometric picture; once it becomes sufficiently familiar it becomes a great source of insight about the properties of quantum operations on a single qubit.

3.2 Freedom in the operator-sum representation

Consider quantum operations $\mathcal{E}$ and $\mathcal{F}$ acting on a single qubit with the operator-sum representations,

$$\begin{aligned} \mathcal{E}(\rho) &= \frac{\rho}{2} + \frac{Z\rho Z}{2} \qquad (3.61) \\ \mathcal{F}(\rho) &= |0\rangle\langle 0|\rho|0\rangle\langle 0| + |1\rangle\langle 1|\rho|1\rangle\langle 1|. \qquad (3.62) \end{aligned}$$

What is interesting is that these two quantum operations are actually the same quantum operation. To see this, note that $|0\rangle\langle 0| = (I + Z)/2$ and $|1\rangle\langle 1| = (I - Z)/2$. Thus

$$\begin{aligned} \mathcal{F}(\rho) &= \frac{(I+Z)\rho(I+Z) + (I-Z)\rho(I-Z)}{4} && (3.63) \\ &= \frac{I\rho I + Z\rho Z}{2} && (3.64) \\ &= \mathcal{E}(\rho). && (3.65) \end{aligned}$$

This freedom in the representation is very interesting. Suppose we flipped a fair coin, and, depending on the outcome of the coin toss, applied either the unitary operator I or Z to the quantum system. This process corresponds to the first operator-sum representation for $\mathcal{E}$. The second operator-sum representation for $\mathcal{E}$ (labeled $\mathcal{F}$ above) corresponds to performing a projective measurement in the $\{|0\rangle, |1\rangle\}$ basis, with the outcome of the measurement unknown. These two apparently very different physical processes give rise to exactly the same system dynamics.

In this section we study in more detail the question of when two sets of operators give rise to the same quantum operation. Understanding this question is important for at least two different reasons. First, from a physical point of view, understanding the freedom in the representation gives us more insight into how different physical processes can give rise to the same system dynamics. Second, in later chapters we will have occasion to use the characterization we find to simplify certain constructions. In particular, it will simplify some of the constructions involving quantum error correction.

To begin, we actually need to answer a different question. Suppose $|\psi_i\rangle$ is a set of states. We say the set $|\psi_i\rangle$ *generates* the operator $A \equiv \sum_i |\psi_i\rangle\langle\psi_i|$. When do two sets of states, $|\psi_i\rangle$ and $|\phi_j\rangle$ generate the same operator A? It turns out that the answer to this question has a surprising number of interesting and useful consequences, amongst which is the solution to our problem of determining the freedom in the operator-sum representation.

Theorem 2 *The sets $|\psi_i\rangle$ and $|\phi_j\rangle$ generate the same operator if and only if*

$$|\psi_i\rangle = \sum_j u_{ij} |\phi_j\rangle, \tag{3.66}$$

where u_{ij} is a unitary matrix of complex numbers, and we "pad" whichever set of states $|\psi_i\rangle$ or $|\phi_j\rangle$ is smaller with additional states 0 so that the two sets have the same number of elements.

As an example of the theorem, suppose we have

$$\rho = \frac{3}{4}|0\rangle\langle 0| + \frac{1}{4}|1\rangle\langle 1|. \tag{3.67}$$

Let

$$\begin{aligned} |a\rangle &\equiv \frac{1}{\sqrt{2}}\sqrt{\frac{3}{4}}|0\rangle + \frac{1}{\sqrt{2}}\sqrt{\frac{1}{4}}|1\rangle && (3.68) \\ |b\rangle &\equiv \frac{1}{\sqrt{2}}\sqrt{\frac{3}{4}}|0\rangle - \frac{1}{\sqrt{2}}\sqrt{\frac{1}{4}}|1\rangle. && (3.69) \end{aligned}$$

Then it is easily checked that $\rho = |a\rangle\langle a| + |b\rangle\langle b|$.

Proof

Suppose $|\psi_i\rangle = \sum_j u_{ij}|\phi_j\rangle$ for some unitary u_{ij}. Then

$$\begin{aligned}
\sum_i |\psi_i\rangle\langle\psi_i| &= \sum_{ijk} u_{ij}u_{ik}^*|\phi_j\rangle\langle\phi_k| && (3.70)\\
&= \sum_{jk}\left(\sum_i u_{ki}^\dagger u_{ij}\right)|\phi_j\rangle\langle\phi_k| && (3.71)\\
&= \sum_{jk}\delta_{kj}|\phi_j\rangle\langle\phi_k| && (3.72)\\
&= \sum_j |\phi_j\rangle\langle\phi_j|, && (3.73)
\end{aligned}$$

which shows that $|\psi_i\rangle$ and $|\phi_j\rangle$ generate the same operator.

Conversely, suppose

$$A = \sum_i |\psi_i\rangle\langle\psi_i| = \sum_j |\phi_j\rangle\langle\phi_j|. \quad (3.74)$$

A little thought shows that for this equation to hold each $|\psi_i\rangle$ can be expressed as a linear combination of the $|\phi_j\rangle$, $|\psi_i\rangle = \sum_j c_{ij}|\phi_j\rangle$. Thus

$$\sum_j |\phi_j\rangle\langle\phi_j| = \sum_{j_1 j_2}\left(\sum_i c_{ij_1}c_{ij_2}^*\right)|\phi_{j_1}\rangle\langle\phi_{j_2}|, \quad (3.75)$$

from which we see that c is unitary, as required.

QED

This result allows us to characterize the freedom in operator-sum representations. Suppose E_j and F_k are two sets of operators, both giving rise to the same quantum operation, $\sum_j E_j A E_j^\dagger = \sum_k F_k A F_k^\dagger$ for all A. Define

$$\begin{aligned}
|e_j\rangle &\equiv \sum_i |i_R\rangle\left(E_j|i_Q\rangle\right) && (3.76)\\
|f_k\rangle &\equiv \sum_i |i_R\rangle\left(F_k|i_Q\rangle\right). && (3.77)
\end{aligned}$$

Recall, the earlier definition of σ, equation (3.26), from which it follows that $\sigma = \sum_j |e_j\rangle\langle e_j| = \sum_k |f_k\rangle\langle f_k|$, and thus there exists unitary u_{jk} such that

$$|e_j\rangle = \sum_k u_{jk}|f_k\rangle. \quad (3.78)$$

But for arbitrary $|\psi\rangle$ we have

$$\begin{aligned}
E_j|\psi\rangle &= \langle\tilde{\psi}|e_j\rangle && (3.79)\\
&= \sum_k u_{jk}\langle\tilde{\psi}|f_k\rangle && (3.80)\\
&= \sum_k u_{jk}F_k|\psi\rangle. && (3.81)
\end{aligned}$$

Thus

$$E_j = \sum_k u_{jk} F_k. \tag{3.82}$$

Conversely, supposing E_j and F_k are related by a unitary transformation of the form $E_j = \sum_{jk} u_{jk} F_k$, simple algebra shows that the quantum operation generated by the operators E_j is the same as the quantum operation generated by the operators F_k.

Summarizing, we have shown that a quantum operation $\mathcal{E}$ is generated in the operator-sum representation by two sets of operators E_j and F_k if and only if there exists a unitary matrix of complex numbers u_{jk} such that

$$E_j = \sum_k u_{jk} F_k, \tag{3.83}$$

where it may be necessary to "pad" the shorter set of operators with zero operators to ensure that the matrix u is square.

This result is surprisingly useful. We will use it, for example, in our study of quantum error correction in Chapter 9. In that Chapter we will see that certain sets operators in the operator sum representation give more useful information about the quantum error correction process, and it behooves us to study quantum error correction from that point of view. As usual, having multiple ways of understanding a process gives us much more insight into what is going on.

3.3 Teleportation as a quantum operation

Let's switch gears, and move away from abstract generalities into a more specific scenario: quantum teleportation. As discussed in section 2.3, quantum teleportation allows us to transmit an unknown quantum state from one location to another using preshared entanglement and classical communication. In this section we show how quantum teleportation can be understood within the quantum operations formalism. This, in turn, allows us to relate quantum teleportation to quantum error correction. The work in this section is based upon a collaboration with Caves [135]. Some of the ideas were arrived at independently about the same time by Bennett, DiVincenzo, Smolin and Wootters [22]. I would especially like to thank Chris Fuchs, who got this work started by suggesting that it might be valuable to try to understand quantum teleportation in terms of reversible measurements.

Recall that teleportation involves a sender, Alice, and a receiver, Bob. Suppose Alice has possession of an *input system*, which we label 1, in an unknown input state $\tilde{\rho}^1$. To avoid confusion, we use a superscript to denote the appropriate state space for a vector or an operator; the reason for the tilde becomes clear shortly. Alice might also have access to another system, which we label 2. Bob has access to the *target system*, which we label 3. Systems 2 and 3 are assumed to be prepared initially in some standard state σ^{23}, which is assumed to be uncorrelated with $\tilde{\rho}^1$; that is, the initial state of the composite system consisting of 1, 2, and 3 is

$$\tilde{\rho}^1 \otimes \sigma^{23} \ . \tag{3.84}$$

The case where Bob has access to an additional system, labeled 4, is discussed briefly later in this section.

We assume that systems 1 and 3 are identical and thus have the same state space. This means that there is a one-to-one linear map from the state space of 3 onto the state space of 1. Though this map is not unique, we choose a particular one, thereby setting up a one-to-one correspondence

between vectors in the state space of 3 and vectors in the state space of 1. We denote this one-to-one correspondence by

$$|\psi^3\rangle \leftrightarrow |\tilde{\psi}^1\rangle \ . \tag{3.85}$$

The one-to-one correspondence between vectors induces a one-to-one correspondence between operators on 3 and operators on 1, which we denote by $A^3 \leftrightarrow \tilde{A}^1$. This correspondence is given by linearly extending the map $|\psi^3\rangle\langle\phi^3| \leftrightarrow |\tilde{\psi}^1\rangle\langle\tilde{\phi}^1|$ to all operators on systems 3 and 1. In particular, for each state $\tilde{\rho}^1$ of the input system, there is a unique counterpart state ρ^3 of the target system.

The choice of a correspondence between the state spaces of 1 and 3 is *physically* motivated: the correspondence defines what it means to transport a system unchanged from the location of system 1 to the location of system 3. Different procedures for performing this transportation lead to different correspondences. For example, suppose we wish to teleport the state of a spin-$\frac{1}{2}$ particle from Los Alamos to Pasadena. To say what it means to teleport the state requires a correspondence between the state spaces in Los Alamos and Pasadena. We could set up the correspondence by agreeing that the z axis in each location lies along the local acceleration of gravity and the x axis along the local magnetic north or by adopting arbitrary orthogonal axes in the two locations. Ordinarily we assume implicitly such a correspondence, as was done earlier in the Dissertation, and write $\tilde{\rho}^1 = \rho^3 = \rho$. In the present setting it is advantageous to adopt a notation which more explicitly distinguishes between states of system 1 and of system 3.

The correspondence can be extended to a one-to-one correspondence between the joint state space of 2 and 3 and the joint state space of 1 and 2. If $|b^2\rangle|c^3\rangle$ is a product basis for the joint state space of 2 and 3, this one-to-one correspondence is given by

$$|\psi^{23}\rangle = \sum_{b,c} \alpha_{bc}|b^2\rangle|c^3\rangle \leftrightarrow \sum_{b,c} \alpha_{bc}|\tilde{c}^1\rangle|b^2\rangle = |\tilde{\psi}^{12}\rangle \ . \tag{3.86}$$

This correspondence induces a one-to-one correspondence between operators on the joint state space of 2 and 3 and operators on the joint state space of 1 and 2.

The correspondence can be extended further to a one-to-one linear map from the state space of the composite system 1, 2, and 3 onto itself:

$$|\psi^{123}\rangle \leftrightarrow |\tilde{\psi}^{123}\rangle = U_{13}|\psi^{123}\rangle \ . \tag{3.87}$$

This map is accomplished by a unitary operator U_{13}, which acts on product states according to

$$U_{13}|\tilde{a}^1\rangle|b^2\rangle|c^3\rangle = |\tilde{c}^1\rangle|b^2\rangle|a^3\rangle \tag{3.88}$$

and thus is called the "swap" operator because it swaps the states of systems 1 and 3, while leaving system 2 alone. The swap operator clearly satisfies $(U_{13})^2 = I^{123}$, that is, $U_{13}^\dagger = U_{13}$. When extended to operators on the composite system, the correspondence becomes

$$A^{123} \leftrightarrow \tilde{A}^{123} = U_{13}A^{123}U_{13}^\dagger \ . \tag{3.89}$$

Suppose now that Alice performs a measurement on systems 1 and 2. This measurement will be described by a set of quantum operations $\mathcal{E}_m$ such that $\sum_m \mathcal{E}_m$ is a complete quantum operation. We assume that each $\mathcal{E}_m$ has an operator-sum representation generated by operators $\tilde{E}_{mj}^{12}$ on the systems 1 and 2.

If the measurement has outcome m, then the unnormalized state of the target system 3 after the measurement is given by

$$\hat{\rho}_m^3 = \mathrm{tr}_{12}\left(\sum_j (\tilde{E}_{mj}^{12} \otimes I^3)(\tilde{\rho}^1 \otimes \sigma^{23})[(\tilde{E}_{mj}^{12})^\dagger \otimes I^3]\right) . \tag{3.90}$$

where the caret denotes an unnormalized state.

We now show that $\hat{\rho}^3_m$ is related to ρ^3 by a quantum operation, which we denote $\mathcal{T}_m$. We first notice that

$$\tilde{\rho}^1 \otimes \sigma^{23} = U_{13}(\tilde{\sigma}^{12} \otimes \rho^3)U_{13}^\dagger \,, \tag{3.91}$$

where $\tilde{\sigma}^{12}$ is the counterpart of σ^{23}. Substituting this into (3.90) gives

$$\begin{aligned} \hat{\rho}^3_m &= \mathrm{tr}_{12}\Bigg(\sum_j (\tilde{E}^{12}_{mj} \otimes I^3) \\ &\qquad \times [U_{13}(\tilde{\sigma}^{12} \otimes \rho^3)U_{13}^\dagger][(\tilde{E}^{12}_{mj})^\dagger \otimes I^3]\Bigg) \,. \end{aligned} \tag{3.92}$$

The form of this equation allows us to think of $\hat{\rho}^3_m$ as arising from the following process. The composite system begins in the state $\tilde{\sigma}^{12} \otimes \rho^3$, in which the joint system 1 and 2 is in the state $\tilde{\sigma}^{12}$ and system 3 is in the state ρ^3. After the composite system evolves under the unitary swap operator, a measurement is performed on the joint system 1 and 2, and then the joint system 1 and 2 is discarded. This process ought to seem highly familiar – it is the same process we used to generate selective quantum operations earlier in the Chapter! Of course, it does not matter that this sequence of events does not literally occur; what matters is that it is effectively as if this occurred. Next, we'll explicitly complete the construction of the quantum operation $\mathcal{E}_m$. This having been done, the problem of teleportation is for Bob to reverse the quantum operation $\mathcal{E}_m$. If the reversal can be done, then Bob can recover the state ρ^3 from the output state $\hat{\rho}^3_m = \mathcal{E}_m(\rho^3)$ of system 3.

We write

$$\tilde{\sigma}^{12} = \sum_k p_k |\tilde{s}^{12}_k\rangle\langle\tilde{s}^{12}_k| \,, \tag{3.93}$$

where the vectors $|\tilde{s}^{12}_k\rangle$ make up the complete orthonormal set of eigenvectors of $\tilde{\sigma}^{12}$ in the joint space of 1 and 2. Furthermore, we let $\tilde{\Pi}^{12}_l = |\tilde{P}^{12}_l\rangle\langle\tilde{P}^{12}_l|$ be any complete set of orthogonal one-dimensional projectors for the joint system 1 and 2. Performing the partial trace of Eq. (3.92) in the basis $|\tilde{P}^{12}_l\rangle$ yields

$$\begin{aligned} \hat{\rho}^3_m &= \sum_{jkl}\Big(\sqrt{p_k}\langle\tilde{P}^{12}_l|(\tilde{E}^{12}_{mj} \otimes I^3)U_{13}|\tilde{s}^{12}_k\rangle\Big)\,\rho^3 \\ &\qquad \times \Big(\sqrt{p_k}\langle\tilde{s}^{12}_k|U_{13}^\dagger[(\tilde{E}^{12}_{mj})^\dagger \otimes I^3]|\tilde{P}^{12}_l\rangle\Big) \,. \end{aligned} \tag{3.94}$$

Using the single index n to denote the triple (j,k,l) and defining the system 3 operators

$$\begin{aligned} B^3_{mn} &\equiv \sqrt{p_k}\langle\tilde{P}^{12}_l|(\tilde{E}^{12}_{mj} \otimes I^3)U_{13}|\tilde{s}^{12}_k\rangle \\ &= \sqrt{p_k}\langle\tilde{P}^{12}_l|U_{13}(I^1 \otimes E^{23}_{mj})|\tilde{s}^{12}_k\rangle \,, \end{aligned} \tag{3.95}$$

we can write the output state of system 3 as

$$\hat{\rho}^3_m = \sum_n B^3_{mn}\rho^3(B^3_{mn})^\dagger \equiv \mathcal{E}_m(\rho^3) \,. \tag{3.96}$$

As we set out to show, $\hat{\rho}^3_m$ is related to ρ^3 by a quantum operation $\mathcal{E}_m$.

We have shown how to construct a quantum operation explicitly linking the the input to the teleportation process to the output. The exact form of the quantum operation depends upon how the teleportation process is performed. In collaboration with Caves I have used this description elsewhere to obtain necessary and sufficient conditions for teleportation, for a subclass of possible teleportation processes [133]. In the present context, the importance of this discussion is as an example of how the quantum operations formalism may be used to obtain explicit representations for interesting quantum processes. In Chapters 9 and 10 we will study the problem of quantum error correction, which turns out to be closely related to teleportation. The connection is to note that for Bob to complete the teleportation process, he must perform a complete quantum operation $\mathcal{R}_m$ on system 3 such that

$$\mathcal{R}_m\left(\frac{\mathcal{E}_m(\rho^3)}{\mathrm{tr}\big(\mathcal{E}_m(\rho^3)\big)}\right) = \rho^3 \, . \tag{3.97}$$

That is, Bob must be able to reverse the quantum operation $\mathcal{E}_m$, recovering the original state ρ^3. The subject of quantum error correction is actually the study of when such a reversal is possible; thus the connection between quantum teleportation and quantum error correction.

3.4 Quantum process tomography

Suppose an experimentalist wishes to completely characterize the dynamics of a quantum system. For finite dimensional systems we explain in this section how this task can be performed with the aid of the quantum operations formalism, and a process known as *quantum state tomography* [146, 112, 111]. The resulting procedure is called *quantum process tomography*, since it gives a method for completely characterizing a quantum process. The work in this section is based upon work done in collaboration with Chuang [40]. Similar work was done independently by Poyatos, Cirac and Zoller [142] at about the same time. Some of these questions have been considered in a partial manner by other researchers, including Jones [91], Turchette *et al* [176] and Mabuchi [121].

The experimental procedure may be outlined as follows. Suppose the state space of the system has N dimensions; for example, $N = 2$ for a single qubit. N^2 pure quantum states $|\psi_1\rangle\langle\psi_1|, \ldots, |\psi_{N^2}\rangle\langle\psi_{N^2}|$ are prepared, and the output state $\mathcal{E}(|\psi_j\rangle\langle\psi_j|)$ is measured for each input. In general, performing such a measurement is not easy, but in recent years a procedure known as *quantum state tomography* [146, 112, 111] has been developed which enables such measurements to be performed. In principle, the quantum operation $\mathcal{E}$ can now be determined by a linear extension of $\mathcal{E}$ to all states, provided the input operators $|\psi_1\rangle\langle\psi_1|, \ldots, |\psi_{N^2}\rangle\langle\psi_{N^2}|$ form a linearly independent set.

From a purist's point of view, we are done. In practice, of course, we would like to have a way of determining a useful representation of $\mathcal{E}$ from experimentally available data. In this section we give a general description of such a method, and an example of how it may be applied in the single qubit case.

Our goal is to determine a set of operators, E_i, generating an operator-sum representation for $\mathcal{E}$,

$$\mathcal{E}(\rho) = \sum_i E_i \rho E_i^\dagger . \tag{3.98}$$

However, experimental results involve numbers, not operators, which are a theoretical concept. To determine the E_i from measurable parameters, it is convenient to consider an equivalent description

of $\mathcal{E}$ using a *fixed* set of operators $\tilde{E}_i$, which form a basis for the set of operators on the state space, so that

$$E_i = \sum_m e_{im} \tilde{E}_m \tag{3.99}$$

for some set of complex numbers e_{im}. Eq.(3.98) may thus be rewritten as

$$\mathcal{E}(\rho) = \sum_{mn} \tilde{E}_m \rho \tilde{E}_n^\dagger \chi_{mn} \,, \tag{3.100}$$

where $\chi_{mn} \equiv \sum_i e_{im} e_{in}^*$ is a matrix which is positive Hermitian by definition.

In general, χ will contain $N^4 - N^2$ independent real parameters, because a general linear map of N by N complex matrices to N by N matrices is described by N^4 independent parameters, but there are N^2 additional constraints due to the fact that ρ remains Hermitian with trace one; that is, the completeness relation

$$\sum_i E_i^\dagger E_i = I, \tag{3.101}$$

is satisfied, giving N^2 real constraints. Note that the restriction that the map be a quantum operation does not change the counting, since by Choi's results [36], the set of quantum operations is just the positive cone in the real vector space of Hermitian-preserving maps, and the positive cone of a real vector space has the same dimensionality as the underlying vector space. We will show how to determine χ experimentally, and then show how an operator-sum representation of the form (3.98) can be recovered once the χ matrix is known.

Let ρ_j, $1 \le j \le N^2$ be a fixed set of linearly independent basis elements for the space of $N \times N$ matrices. A convenient choice is the set of operators $|n\rangle\langle m|$. Experimentally, the output state $\mathcal{E}(|n\rangle\langle m|)$ may be obtained by preparing the input states $|n\rangle$, $|m\rangle$, $|n_+\rangle = (|n\rangle + |m\rangle)/\sqrt{2}$, and $|n_-\rangle = (|n\rangle + i|m\rangle)/\sqrt{2}$ and forming linear combinations of $\mathcal{E}(|n\rangle\langle n|)$, $\mathcal{E}(|m\rangle\langle m|)$, $\mathcal{E}(|n_+\rangle\langle n_+|)$, and $\mathcal{E}(|n_-\rangle\langle n_-|)$. Thus, it is possible to determine $\mathcal{E}(\rho_j)$ by state tomography, for each ρ_j.

Furthermore, each $\mathcal{E}(\rho_j)$ may be expressed as a linear combination of the basis states,

$$\mathcal{E}(\rho_j) = \sum_k \lambda_{jk} \rho_k \,, \tag{3.102}$$

and since $\mathcal{E}(\rho_j)$ is known, λ_{jk} can be determined by standard linear algebraic algorithms. To proceed, we may write

$$\tilde{E}_m \rho_j \tilde{E}_n^\dagger = \sum_k \beta_{jk}^{mn} \rho_k \,, \tag{3.103}$$

where β_{jk}^{mn} are complex numbers which can be determined by standard algorithms from linear algebra given the $\tilde{E}_m$ operators and the ρ_j operators. Combining the last two expressions we have

$$\sum_k \sum_{mn} \chi_{mn} \beta_{jk}^{mn} \rho_k = \sum_k \lambda_{jk} \rho_k \,. \tag{3.104}$$

From independence of the ρ_k it follows that for each k,

$$\sum_{mn} \beta_{jk}^{mn} \chi_{mn} = \lambda_{jk} \,. \tag{3.105}$$

This relation is a necessary and sufficient condition for the matrix χ to give the correct quantum operation $\mathcal{E}$. One may think of χ and λ as vectors, and β as a $N^4 \times N^4$ matrix with columns indexed by mn, and rows by ij. To show how χ may be obtained, let κ be the generalized inverse for the matrix β, satisfying the relation

$$\beta_{jk}^{mn} = \sum_{st,xy} \beta_{jk}^{st} \kappa_{st}^{xy} \beta_{xy}^{mn} \,. \tag{3.106}$$

Most computer packages for matrix manipulation are capable of finding such generalized inverses. We now prove that χ defined by

$$\chi_{mn} = \sum_{jk} \kappa_{jk}^{mn} \lambda_{jk} \tag{3.107}$$

satisfies the relation (3.105).

The difficulty in verifying that χ defined by (3.107) satisfies (3.105) is that, in general, χ is not uniquely determined by equation (3.105). For convenience we will rewrite these equations in matrix form as

$$\beta \vec{\chi} = \vec{\lambda} \tag{3.108}$$

$$\vec{\chi} \equiv \kappa \vec{\lambda} \,. \tag{3.109}$$

From the construction that led to equation (3.100) we know there exists at least one solution to equation (3.108), which we shall call $\vec{\chi}'$. Thus $\vec{\lambda} = \beta \vec{\chi}'$. The generalized inverse satisfies $\beta \kappa \beta = \beta$. Premultiplying the definition of $\vec{\chi}$ by β gives

$$\beta \vec{\chi} = \beta \kappa \vec{\lambda} \tag{3.110}$$

$$= \beta \kappa \beta \vec{\chi}' \tag{3.111}$$

$$= \beta \vec{\chi}' \tag{3.112}$$

$$= \lambda \,. \tag{3.113}$$

Thus χ defined by (3.109) satisfies the equation (3.108), as we wanted to show.

Having determined χ one immediately obtains the operator-sum representation for $\mathcal{E}$ in the following manner. Let the unitary matrix $U^\dagger$ diagonalize χ,

$$\chi_{mn} = \sum_{xy} U_{mx} d_x \delta_{xy} U_{ny}^* \,. \tag{3.114}$$

From this it can easily be verified that

$$E_i = \sqrt{d_i} \sum_j U_{ji} \tilde{E}_j \tag{3.115}$$

gives an operator-sum representation for the quantum operation $\mathcal{E}$. Our algorithm may thus be summarized as follows: λ is experimentally measured, and given β, determined by a choice of $\tilde{E}$, we find the desired parameters χ which completely describe $\mathcal{E}$, and which determine a set of operators E_i generating an operator-sum representation for $\mathcal{E}$.

3.4.1 One qubit example

The above general method can be simplified in the case of a one qubit operation to provide explicit formulas which may be useful in experimental contexts, such as the teleportation experiment

described in subsection 2.6.2. This simplification is made possible by choosing the fixed operators $\tilde{E}_i$ to have commutation properties which conveniently allow the χ matrix to be determined by straightforward matrix multiplication. In the one qubit case, we use:

$$\tilde{E}_0 = I \tag{3.116}$$
$$\tilde{E}_1 = X \tag{3.117}$$
$$\tilde{E}_2 = -iY \tag{3.118}$$
$$\tilde{E}_3 = Z. \tag{3.119}$$

There are 12 parameters, specified by χ, which determine an arbitrary single qubit quantum operation $\mathcal{E}$. These 12 parameters may be measured using four sets of experiments. As a specific example, suppose the input states $|0\rangle$, $|1\rangle$, $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and $|-\rangle = (|0\rangle + i\,|1\rangle)/\sqrt{2}$ are prepared, and the four matrices

$$\rho'_1 = \mathcal{E}(|0\rangle\langle 0|) \tag{3.120}$$
$$\rho'_4 = \mathcal{E}(|1\rangle\langle 1|) \tag{3.121}$$
$$\rho'_2 = \mathcal{E}(|+\rangle\langle +|) - i\mathcal{E}(|-\rangle\langle -|) - (1-i)(\rho'_1 + \rho'_4)/2 \tag{3.122}$$
$$\rho'_3 = \mathcal{E}(|+\rangle\langle +|) + i\mathcal{E}(|-\rangle\langle -|) - (1+i)(\rho'_1 + \rho'_4)/2 \tag{3.123}$$

are determined using state tomography. These correspond to $\rho'_j = \mathcal{E}(\rho_j)$, where

$$\rho_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \tag{3.124}$$

$\rho_2 = \rho_1 X$, $\rho_3 = X\rho_1$, and $\rho_4 = X\rho_1 X$. From Eq.(3.103) and Eqs.(3.116-3.119) we may determine β, and similarly ρ'_j determines λ. However, due to the particular choice of basis, and the Pauli matrix representation of $\tilde{E}_i$, we may express the β matrix as the Kronecker product $\beta = \Lambda \otimes \Lambda$, where

$$\Lambda = \frac{1}{2}\begin{bmatrix} I & X \\ X & -I \end{bmatrix}, \tag{3.125}$$

so that χ may be expressed conveniently as

$$\chi = \Lambda \begin{bmatrix} \rho'_1 & \rho'_2 \\ \rho'_3 & \rho'_4 \end{bmatrix} \Lambda\,, \tag{3.126}$$

in terms of block matrices.

Consider a one-qubit black box of unknown dynamics $\mathcal{E}_1$. Suppose that the following four density matrices are obtained from experimental measurements, performed according to Eqs.(3.120-3.123):

$$\rho'_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \tag{3.127}$$
$$\rho'_2 = \begin{bmatrix} 0 & \sqrt{1-\gamma} \\ 0 & 0 \end{bmatrix} \tag{3.128}$$
$$\rho'_3 = \begin{bmatrix} 0 & 0 \\ \sqrt{1-\gamma} & 0 \end{bmatrix} \tag{3.129}$$
$$\rho'_4 = \begin{bmatrix} \gamma & 0 \\ 0 & 1-\gamma \end{bmatrix}, \tag{3.130}$$

where γ is a numerical parameter. From a independent study of each of these input-output relations, one could make several important observations: the ground state $|0\rangle$ is left invariant by $\mathcal{E}_1$, the excited state $|1\rangle$ partially decays to the ground state, and superposition states are damped.

However, let us proceed systematically and determine χ using this data. From Eqs.(3.125-3.126), we find the χ matrix for this process to be

$$\chi = \frac{1}{4}\begin{bmatrix} (1+\sqrt{1-\gamma})^2 & 0 & 0 & \gamma \\ 0 & \gamma & -\gamma & 0 \\ 0 & -\gamma & \gamma & 0 \\ \gamma & 0 & 0 & (1-\sqrt{1-\gamma})^2 \end{bmatrix} \tag{3.131}$$

Using Eqs.(3.114-3.115), we then obtain (after a little simplification) the operators E_i which generate the operator-sum representation for this quantum operation,

$$E_0 = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{bmatrix} \tag{3.132}$$

$$E_1 = \begin{bmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{bmatrix}. \tag{3.133}$$

These operators define a well-known process called *amplitude damping*. It can result from a relaxation process with a microscopic interaction Hamiltonian of the form $\mathcal{H}_I = \gamma'(\sigma^- b^\dagger + \sigma^+ b)$, where σ^+ and $b^\dagger$ are system and environment creation operators, and γ is related to γ' and the interaction time. This process is important, for instance, in quantum error correction, where one wishes to reverse the effects of noise, because better codes exist to correct amplitude damping than for general error processes[113].

The dynamics of a two-qubit quantum black box $\mathcal{E}_2$ pose an even greater challenge for our understanding. In this case there are 240 parameters which need to be determined in order to do completely specify the quantum operation acting on the quantum system! This is obviously quite a considerable undertaking, however, as for the single qubit case, it is relatively straightforward to implement a numerical routine which will automate the calculation, provided experimental tomography and state preparation are available in the laboratory. We will not give an example here, as it does not serve the purpose of the present Chapter, referring the reader instead to [40] for more details.

Until now we have been considering complete quantum operations. In a situation where quantum measurements may be involved, the corresponding quantum operations may be incomplete. We now briefly outline how to determine the quantum operation corresponding to each measurement outcome in this instance.

Recall that for each measurement outcome, m, there is associated a quantum operation, $\mathcal{E}_m$. The corresponding state change is given by

$$\rho \to \frac{\mathcal{E}_m(\rho)}{\mathrm{tr}(\mathcal{E}_m(\rho))}, \tag{3.134}$$

where the probability of the measurement outcome occurring is $p_m = \mathrm{tr}(\mathcal{E}_m(\rho))$. Note that this mapping is nonlinear, because of this renormalization factor, so the earlier methods do not apply.

Despite the possible nonlinearity, the procedure we have described may be adapted in a straightforward manner to evaluate the quantum operations describing a measurement. To determine $\mathcal{E}_m$ we proceed exactly as before, except now we must perform the measurement a large enough number of times that the probability p_m can be reliably estimated, for example, by using the

frequency of occurrence of outcome m. This must be done for each input ρ_j which is to be used for the tomography procedure. Note that standard statistical tools may be used to estimate the accuracy with which the probability p_m has been determined. Once the p_m have all been estimated to some desired accuracy, ρ'_j is determined using state tomography, allowing us to obtain

$$\mathcal{E}_m(\rho_j) = p_m \rho'_j, \tag{3.135}$$

for each input ρ_j which we prepare, since each term on the right hand side is known. Now we proceed exactly as before to evaluate the quantum operation $\mathcal{E}_i$.

Summing up, we have shown how a useful representation for the dynamics of a quantum system may be experimentally determined using a systematic procedure. This procedure of *quantum process tomography* is analogous to the *system identification* step [116] performed in classical control theory. Quantum process tomography opens the way for robust experimental determination of a wide variety of interesting quantities associated to noisy quantum processes. As such, I expect it will eventually become an indispensable tool in the experimental study of quantum information processing.

3.5 The POVM formalism

One of the main uses of the quantum operations formalism is to describe the effects of measurement. Quantum operations can be used to describe both the probability of getting a particular outcome from a measurement on a quantum system, and also the state change in the system effected by the measurement.

In many cases, though, the state change in the system being measured is not particularly interesting, since the system itself is discarded after the measurement is performed. For example, this is the case for photons detected by a photodetector, which destroys the photon.

What is still interesting in these examples is the probabilities of different measurement outcomes. It turns out that the quantum operations formalism simplifies rather nicely if one is only interested in the probabilities of different measurement outcomes, and not also the corresponding state changes. This simplified formalism has become known for historical reasons as the *Positive Operator Valued Measure* formalism, or POVM formalism for short.

You may ask why we should bother studying a formalism which is a special case of a more general formalism. The reason is that it sometimes simplify matters to consider a problem from the point of view of POVMs. New sources of intuition in quantum information are to be valued, and the simplicity of the POVM formalism is one such source of intuition.

Suppose we consider a set of Hermitian operators M_m indexed by an index which we denote m, satisfying the conditions

$$M_m \geq 0 \tag{3.136}$$

$$\sum_m M_m = I. \tag{3.137}$$

Consider now a measurement described by quantum operations $\mathcal{E}_m$ defined by the equations $\mathcal{E}_m(\rho) = \sqrt{M_m}\rho\sqrt{M_m}$. Notice that the quantum operation $\sum_m \mathcal{E}_m$ is a complete quantum operation by equation (3.137), and that the probability of outcome m occurring is given by

$$p(m) = \mathrm{tr}(\sqrt{M_m}\rho\sqrt{M_m}) = \mathrm{tr}(M_m\rho). \tag{3.138}$$

Thus, given a set of operators M_m satisfying the conditions (3.136) and (3.137), it is possible to find a measurement model such that equation (3.138) correctly gives the probability of the measurement outcome m.

Conversely, suppose a measurement is taking place, which is described by quantum operations $\mathcal{E}_m$ associated to the measurement outcomes m. Let E_{mi} be a set of operators generating the quantum operation $\mathcal{E}_m$. Define $M_m \equiv \sum_i E_{mi}^\dagger E_{mi}$. Note that $M_m \geq 0$, and

$$\begin{aligned} \mathrm{tr}(\mathcal{E}_m(\rho)) &= \sum_i \mathrm{tr}(E_{mi}\rho E_{mi}^\dagger) && (3.139) \\ &= \mathrm{tr}(M_m\rho), && (3.140) \end{aligned}$$

so $\mathrm{tr}(M_m\rho)$ gives the probability of outcome m occurring in the measurement. The completeness relation $\sum_{mi} E_{mi}^\dagger E_{mi} = I$ is true if and only if $\sum_m M_m = I$.

These two results suggest the following formal definition. A *POVM* consists of a set of operators M_m satisfying the two conditions:

1. (*Positivity*)

$$M_m \geq 0. \tag{3.141}$$

2. (*Completeness*)

$$\sum_m M_m = I. \tag{3.142}$$

A POVM describes the probabilities of the measurement outcomes via the rule

$$p(m) = \mathrm{tr}(M_m\rho). \tag{3.143}$$

These three equations – the positivity requirement, completeness, and the probability rule – completely summarize the POVM formalism. Our results imply that any description of a quantum measurement in terms of quantum operations gives rise to a unique POVM describing the measurement statistics for that measurement. We have also shown that given any POVM, there exists a measurement model whose statistics agree with those predicted by the POVM.

3.6 Beyond quantum operations?

Are there interesting quantum systems whose dynamics are not described by quantum operations? In this section we give a very brief discussion of this question. A more detailed discussion of some of these issues has been provided by Royer [147]. In this section we will construct an artificial example of a system whose evolution is not described by a quantum operation, and try to understand the circumstances under which this is likely to occur.

Suppose a single qubit is prepared in some unknown quantum state, which we denote ρ. The preparation of this qubit involves certain procedures to be carried out in the laboratory in which the qubit is prepared. Suppose that amongst the laboratory degrees of freedom is a single qubit which, as a side effect of the state preparation procedure, is left in the state $|0\rangle$ if ρ is a state on the bottom half of the Bloch sphere, and is left in the state $|1\rangle$ if ρ is a state on the top half of the Bloch sphere. That is, the state of the system after preparation is

$$\rho \otimes |0\rangle\langle 0| \otimes \text{other degrees of freedom} \tag{3.144}$$

if ρ is a state on the bottom half of the Bloch sphere, and

$$\rho \otimes |1\rangle\langle 1| \otimes \text{other degrees of freedom} \tag{3.145}$$

if ρ is a state on the top half of the Bloch sphere.

Once the state preparation is done, the system begins to interact with the environment, in this case all the laboratory degrees of freedom. Suppose the interaction is such that a controlled not is performed between the principal system and the extra qubit in the laboratory system. Thus, if the system's Bloch vector was initially in the bottom half of the Bloch sphere it is left invariant by the process, while if it was initially in the top half of the Bloch sphere it is rotated into the bottom half of the Bloch sphere.

Obviously, this process is not an affine map acting on the Bloch sphere, and therefore, by the results of subsection 3.1.1, it *can not be a quantum operation.* The lesson to be learned from this discussion is that *a quantum system which interacts with the degrees of freedom used to prepare that system after the preparation is complete will in general suffer a dynamics which is not adequately described within the quantum operations formalism.* This is an important conclusion to have reached, as it indicates that there are physically reasonable circumstances under which the quantum operations formalism may not adequately describe the processes taking place in a quantum system. This should be kept in mind, for example, in applications of the quantum process tomography procedure discussed in the previous section.

For the remainder of this Dissertation we will, however, work within the quantum operations formalism. It provides a powerful, and reasonably general tool for describing the dynamics experienced by quantum systems. Most of all, it provides a means by which concrete progress can be made on problems related to quantum information processing. It is an interesting problem for further research to study quantum information processing beyond the quantum operations formalism.

Summary of Chapter 3: Quantum operations

- **Axioms for complete quantum operations:** Linear maps on density operators which preserve trace, and preserve positivity of density operators, even when extended in a natural way to larger systems.
- **Operator-sum representation for a quantum operation:**

$$\mathcal{E}(\rho) = \sum_i E_i \rho E_i^\dagger.$$

  The quantum operations generated by operators E_i and F_j in the operator-sum representation are the same if and only if there exists a unitary matrix u_{ij} such that $E_i = \sum_j u_{ij} F_j$. It may be necessary to append 0 operators so that both sets of operators have the same number of elements.
- **Environmental models for quantum operations:** A complete quantum operation can always be regarded as arising from the unitary interaction of a system with an initially uncorrelated environment, and vice versa. Incomplete quantum operations may be treated similarly, except an additional projective measurement is performed on the composite of system and environment, with the different outcomes corresponding to different incomplete quantum operations.
- **Quantum teleportation:** The input and output states to the quantum teleportation procedure are related by a set of quantum operations. The problem of teleportation is to reverse or error correct those quantum operations.
- **Quantum process tomography:** A procedure used to completely characterize the dynamics of a quantum system in the laboratory.

Chapter 4

Entropy and information

Entropy is a key concept of quantum information theory. It measures how much uncertainty there is in the state of a physical system. In this Chapter we review the basic definitions and properties of entropy in both classical and quantum information theory. In places the Chapter contains rather detailed and lengthy mathematical arguments; upon a first read, these sections may be read lightly, and returned to later for reference purposes.

4.1 Shannon entropy

The key concept of classical information theory is the *Shannon entropy*. Suppose we learn the value of a random variable, X. The Shannon entropy associated with X quantifies how much information we gain, on average, when we learn the value of X. An alternative view is that the entropy of X measures the amount of *uncertainty* about X before we learn the value of X. These two views are complementary: we can view the entropy either as a measure of uncertainty *before* we learn the value of X, or as a measure of how much information we have gained *after* we learn the value of X.

The entropy of a random variable is completely determined by the probabilities of the different possible values that random variable takes. For that reason, we will often write the entropy as a function of a probability distribution, $p_1, \ldots, p_n$. The *Shannon entropy* associated with that probability distribution is defined by

$$H(X) \equiv H(p_1, \ldots, p_n) \equiv -\sum_i p_i \log p_i. \tag{4.1}$$

We will justify this definition shortly. Note that in the definition – and throughout the Dissertation, unless otherwise noted – logarithms indicated by log are taken to base two, while ln indicates a natural logarithm. The reader may wonder what happens when $p_i = 0$, since $\log 0$ is undefined. Intuitively, an event which can never occur should not contribute to the entropy, so by convention we agree that $0 \log 0 \equiv 0$. More formally, note that $\lim_{x \rightarrow 0} x \log x = 0$, which provides further support for our intuition, and thus our convention.

Why is the entropy defined in this way? In the pedagogical literature, it is common to give an axiomatic characterization of the entropy, based upon certain intuitive properties we would expect a measure of information to possess (see, for example [54, 55] for excellent pedagogical introductions to information theory which contain such characterizations). These axioms are then used to deduce the above formula for entropy. While appealing, there is a better reason than axiomatics for choosing

this definition for entropy. The better reason for this definition of entropy is that it can be used to quantify the resources needed to store information. More concretely, suppose there is some *source* (perhaps a radio antenna) which is producing information, say in the form of a bit string. Let's consider a very simple model for a source: we describe it as producing a string $X_1, X_2, \ldots$ of independent, identically distributed random variables. Most real sources don't behave quite that way, but often it's a good approximation. Shannon asked what minimal physical resources are required to store the information being produced by the source, in such a way that at a later time the original source information can be reconstructed [160, 162]? The answer to this question turns out to be the entropy, that is, $H(X_1)$ bits are required per source symbol. This result is known as *Shannon's noiseless coding theorem*, and we will prove both classical and quantum versions of it in Chapter 7.

More abstractly, this motivation for the definition of entropy expresses one of the key philosophies of information theory, both quantum and classical: *fundamental measures of information arise as the answers to fundamental questions about the quantity of physical resources required to solve some information processing problem.*

4.2 Basic properties of entropy

4.2.1 The binary entropy

The entropy of a two outcome random variable is so useful that we will give it a special name, the *binary entropy function*, defined as

$$H_{\text{bin}}(p) = -p \log p - (1-p) \log(1-p), \tag{4.2}$$

where p and $1-p$ are the probabilities of the two outcomes. Where context makes the meaning clear we will write $H(p)$ rather than $H_{\text{bin}}(p)$. Note again that logarithms will be taken to be base two, unless otherwise stated. The binary entropy function is plotted in figure 4.1. Notice that $H(p) = H(1-p)$ and that $H(p)$ attains its maximum value of 1 at $p = 1/2$.

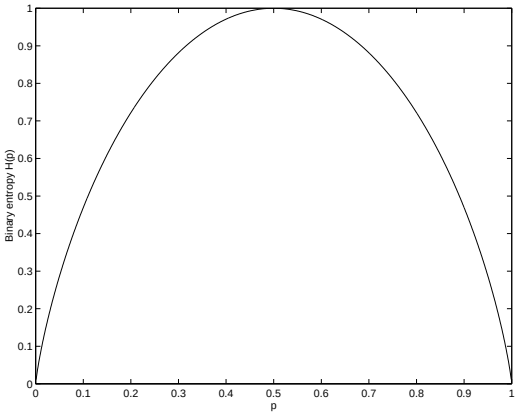

Figure 4.1: Binary entropy function $H(p)$.

4.2.2 The relative entropy

There is a very useful entropy-like measure of the closeness of two probability distributions, $p(x)$ and $q(x)$, over the same index set, x. Suppose $p(x)$ and $q(x)$ are two probability distributions on same index set, x. Define the *relative entropy* of $p(x)$ to $q(x)$ by

$$H(p(x)\|q(x)) \equiv \sum_x p(x) \log \frac{p(x)}{q(x)} \equiv -H(X) - \sum_x p(x) \log q(x). \tag{4.3}$$

We define $-0 \log 0 \equiv 0$ and $-p(x) \log 0 \equiv \infty$ if $p(x) > 0$.

It is probably not immediately obvious what the relative entropy is good for, or even why it is a good measure of distance between two distributions. The following theorem gives some motivation for why it is regarded as being like a distance measure.

Theorem 3 *[54] The relative entropy is non-negative, $H(p(x)\|q(x)) \geq 0$, with equality if and only if $p(x) = q(x)$ for all x.*

As exemplified here, many of the references for elementary results in this Chapter will be to the excellent text of Cover and Thomas [54] or the review paper of Wehrl [182], to which you may refer for historical details.

Proof

A very useful inequality in information theory is $\log x \ln 2 = \ln x \leq x - 1$, for all positive x, with equality if and only if $x = 1$. Here we need to rearrange the result slightly, to $-\log x \ln 2 \geq 1 - x$, and then note that

$$\begin{aligned} H(p(x)\|q(x)) &= -\sum_x p(x) \log \frac{q(x)}{p(x)} && (4.4)\\ &\geq \frac{1}{\ln 2} \sum_x p(x) \left(\frac{q(x)}{p(x)} - 1 \right) && (4.5)\\ &= \frac{1}{\ln 2} \sum_x (q(x) - p(x)) && (4.6)\\ &= \frac{1}{\ln 2}(1 - 1) = 0, && (4.7) \end{aligned}$$

which is the desired inequality. The equality conditions are easily deduced by noting that equality occurs in the second line if and only if $q(x)/p(x) = 1$ for all x, that is, the distributions are identical.

QED

The relative entropy is usually useful, not in itself, but because other entropic quantities can be regarded as special cases of the relative entropy. Theorems about the relative entropy then give as special cases theorems about other entropic quantities. For example, we can use the non-negativity of the relative entropy to prove the following fundamental fact about entropies. Suppose $p(x)$ is a probability distribution for X, over d outcomes. Let $q(x) \equiv 1/d$ be the uniform probability distribution over those outcomes. Then

$$H(p(x)\|q(x)) = H(p(x)\|1/d) = -H(X) - \sum_x p(x) \log(1/d) = \log d - H(X). \tag{4.8}$$

From the non-negativity of the relative entropy we see that $\log d - H(X) \geq 0$, with equality if and only if X is uniformly distributed. This is an elementary fact, but so important that we restate it formally as a theorem.

Theorem 4 *[54] Suppose $p(x)$ is a probability distribution for X, on d outcomes. Then $H(X) \leq \log d$, with equality if and only if $p(x)$ is uniformly distributed.*

We will use this technique – finding expressions for entropic quantities in terms of the relative entropy – often in the study of both classical and quantum entropies. As another example, it is easily verified that $H(p(x,y)\|p(x)p(y)) = H(p(x)) + H(p(y)) - H(p(x,y))$. From this observation and the non-negativity of the relative entropy, we see that $H(X,Y) \leq H(X) + H(Y)$, with equality if and only if X and Y are independent random variables.

4.2.3 Mutual information and conditional entropy

Suppose X and Y are two random variables. How is the information content of X related to the information content of Y? In this subsection we introduce two concepts – the *conditional entropy* and *the mutual information* – which help answer this question. The definitions of these concepts which we give are rather formal, and at times the reader may be confused as to why a particular quantity – say, the conditional entropy – is to be interpreted in the way we indicate. Keep in mind that the ultimate justification for these definitions is that they answer resource questions, which will become clearer in later Chapters. The interpretation given to the quantities depends on the nature of the resource question being answered.

We already met the *joint entropy* of a pair of random variables implicitly in the last subsection. For clarity, we now make this definition formal. The *joint entropy* of X and Y is defined in the obvious way,

$$H(X,Y) \equiv -\sum_{x,y} p(x,y) \log p(x,y), \tag{4.9}$$

and may be extended in the obvious way to any vector of random variables. The joint entropy measures our total uncertainty about the pair (X,Y). Suppose we know the value of Y, so we have acquired $H(Y)$ bits of information about the pair, (X,Y). The remaining uncertainty about the pair (X,Y), is associated with our remaining lack of knowledge about X, even given that we know Y. The *entropy of X conditional on knowing Y* is therefore defined by

$$H(X|Y) \equiv H(X,Y) - H(Y). \tag{4.10}$$

The conditional entropy is a measure of how uncertain we are, on average, about the value of X, given that we know the value of Y.

A second quantity, the *mutual information content of X and Y*, measures how much information X and Y have in common. Suppose we add the information content of X, $H(X)$, to the information content of Y. Then all the information in the pair (X,Y) will have been counted at least once in the sum. Information which is common to X and Y will have been counted twice in this sum, while information which is not common will have been counted only once. Subtracting off the joint information of (X,Y), $H(X,Y)$, we obtain the common or *mutual information* of X and Y:

$$H(X:Y) \equiv H(X) + H(Y) - H(X,Y). \tag{4.11}$$

Notice the useful equality $H(X:Y) = H(X) - H(X|Y)$ relating the conditional entropy and mutual information.

To get some feeling for how the Shannon entropy behaves, we will prove some simple relationships between the different entropies.

Theorem 5 ***(Basic properties of entropy)*** *[54]*

1. $H(X,Y) = H(Y,X)$, $H(X:Y) = H(Y:X)$.
2. $H(Y|X) \geq 0$ *and thus* $H(X:Y) \leq H(Y)$, *with equality if and only if* Y *is a function of* X, $Y = f(X)$.
3. $H(X) \leq H(X,Y)$, *with equality if and only if* Y *is a function of* X.
4. **Subadditivity:** $H(X,Y) \leq H(X) + H(Y)$ *with equality if and only if* X *and* Y *are independent random variables.*
5. $H(Y|X) \leq H(Y)$ *and thus* $H(X:Y) \geq 0$, *with equality in each if and only if* X *and* Y *are independent random variables.*
6. **Strong subadditivity:** $H(X,Y,Z) + H(Y) \leq H(X,Y) + H(Y,Z)$.

Proof

1. Obvious from the relevant definitions.
2. Since $p(x,y) = p(x)p(y|x)$ we have

$$\begin{aligned} H(X,Y) &= -\sum_{xy} p(x,y) \log p(x)p(y|x) && (4.12) \\ &= -\sum_{x} p(x) \log p(x) - \sum_{xy} p(x,y) \log p(y|x) && (4.13) \\ &= H(X) - \sum_{xy} p(x,y) \log p(y|x). && (4.14) \end{aligned}$$

   Thus $H(Y|X) = -\sum_{xy} p(x,y) \log p(y|x)$. But $-\log p(y|x) \geq 0$, so $H(Y|X) \geq 0$ with equality if and only if Y is a deterministic function of X.
3. Follows from the previous result.
4. To prove subadditivity and, later, strong subadditivity we use the fact that $\ln x \leq x - 1$ for all positive x, with equality if and only if $x = 1$. This fact is easily proved using calculus. We find that

$$\begin{aligned} \sum_{x,y} p(x,y) \ln \frac{p(x)p(y)}{p(x,y)} &\leq \sum_{x,y} p(x,y) \left(\frac{p(x)p(y)}{p(x,y)} - 1 \right) && (4.15) \\ &= \sum_{x,y} p(x)p(y) - p(x,y) = 1 - 1 = 0. && (4.16) \end{aligned}$$

   Subadditivity may easily be recovered by multiplying by a constant (to change the base of the logarithm to base 2), and rearranging the expression. Notice that equality is achieved if and only if $p(x,y) = p(x)p(y)$ for all x and y. That is, the subadditivity inequality is saturated if and only if X and Y are independent.
5. Follows from subadditivity and the relevant definitions.

6. Strong subadditivity of Shannon entropy follows from the same technique as used to prove subadditivity; the difficulty level is about the same as that proof. Interestingly, while carrying out the proof one notes that the equality conditions for strong subadditivity are that $Z \rightarrow Y \rightarrow X$ forms a Markov chain.

QED

The various relationships between entropies may mostly be deduced from the "entropy Venn diagram" shown in figure 4.2. These figures are not completely reliable as a guide to the properties of entropy, but they are a useful mnemonic for remembering the various definitions and properties of entropy.

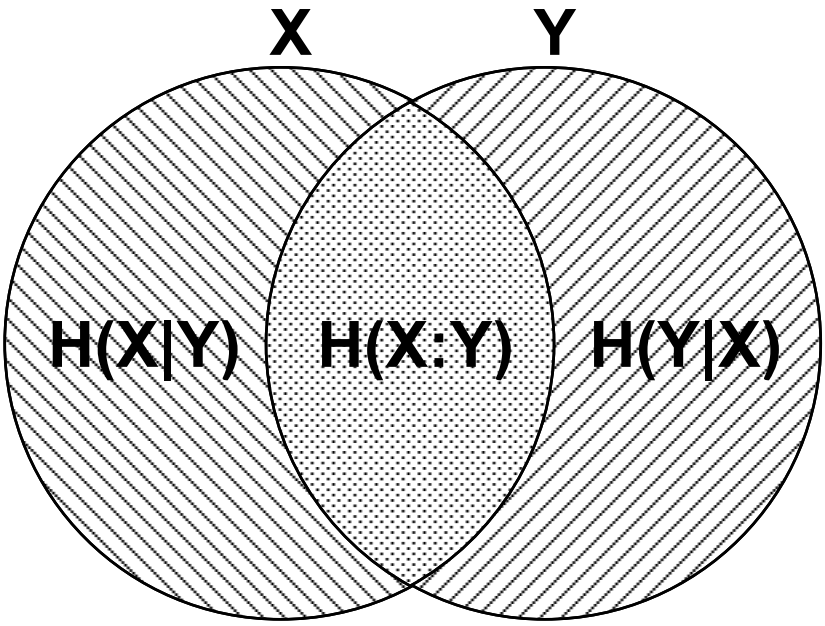

Figure 4.2: Relationships between different entropies.

Intuitively, we expect that the uncertainty about X, given that we know the value of Y and Z, is less than our uncertainty about X, given that we only know Y. More formally,

Theorem 6 ***(Conditioning reduces entropy)*** *[54]*

$$H(X|Y,Z) \leq H(X|Y). \tag{4.17}$$

Proof

Inserting the relevant definitions, the result is equivalent to

$$H(X,Y,Z) - H(Y,Z) \leq H(X,Y) - H(Y), \tag{4.18}$$

which is a rearranged version of the strong subadditivity inequality proved earlier.

QED

The next result gives a simple, useful formula for the conditional entropy.

Theorem 7 ***(Chaining for conditional entropies)*** *[54]*
Let $X_1, \ldots, X_n$ and Y be any set of random variables. Then

$$H(X_1, \ldots, X_n|Y) = \sum_{i=1}^{n} H(X_i|Y, X_1, \ldots, X_{i-1}). \tag{4.19}$$

Proof
We prove the result for $n = 2$, and then induct on n. Using only the definitions and some simple algebra we have

$$\begin{aligned} H(X_1, X_2|Y) &= H(X_1, X_2, Y) - H(Y) && (4.20)\\ &= H(X_1, X_2, Y) - H(X_1, Y) + H(X_1, Y) - H(Y) && (4.21)\\ &= H(X_2|Y, X_1) + H(X_1|Y), && (4.22)\end{aligned}$$

which establishes the result for $n = 2$. Now we assume the result for general n, and show the result holds for $n + 1$. Using the already established $n = 2$ case, we have

$$H(X_1, \ldots, X_{n+1}|Y) = H(X_2, \ldots, X_n|Y, X_1) + H(X_1|Y). \tag{4.23}$$

Applying the inductive hypothesis to the first term on the right hand side gives

$$\begin{aligned} H(X_1, \ldots, X_{n+1}|Y) &= \sum_{i=2}^{n+1} H(X_i|Y, X_1, \ldots, X_{i-1}) + H(X_1|Y) && (4.24)\\ &= \sum_{i=1}^{n+1} H(X_i|Y, X_1, \ldots, X_{i-1}), && (4.25)\end{aligned}$$

so the induction goes through.
QED
Finally, we conclude with a note that will be of interest in Chapter 10 on the quantum channel capacity. In that Chapter we will be much interested in the subadditivity properties of quantities like the mutual information. We now note that the mutual information is not generally subadditive in either or both entries. For instance, let X and Y be independent identically distributed random variables taking the values 0 or 1 with probability 1/2. Let $Z \equiv X + Y$, where the addition is done modulo two. Then it is easy to see that

$$1 = H(X, Y : Z) \not\leq H(X : Z) + H(Y : Z) = 0 + 0. \tag{4.26}$$

Neither is the mutual information superadditive. For example, suppose $X_1 = X_2 = Y_1 = Y_2$, and X_1 is chosen to have the values 0 or 1 with respective probabilities of one half. Then

$$1 = H(X_1, X_2 : Y_1, Y_2) < H(X_1 : Y_1) + H(X_2 : Y_2) = 1 + 1 = 2. \tag{4.27}$$

4.2.4 The data processing inequality

In many applications of interest we perform computations on the information we have available, but that information is imperfect, as it has been subjected to noise before it becomes available to us. A basic inequality of information theory, the *data processing inequality*, states that the information we have available about a source of information can only *decrease with time*: once information has been lost, it is gone forever. Making this statement more precise is the goal of this subsection.

The intuitive notion of *information processing* is captured in the idea of a Markov chain of random variables. Formally, a *Markov chain* is a sequence $X_1, X_2, \ldots$ of random variables such that X_{n+1} is independent of $X_1, \ldots, X_{n-1}$, given X_n. More formally,

$$p(X_{n+1} = x_{n+1}|X_n = x_n, \ldots, X_1 = x_1) = p(X_{n+1} = x_{n+1}|X_n = x_n). \tag{4.28}$$

When is a Markov chain losing information? The following *data processing inequality* gives an information-theoretic way of answering this question.

Theorem 8 (Data processing inequality) *[54]*
Suppose $X \rightarrow Y \rightarrow Z$ is a Markov chain. Then

$$H(X) \geq H(X : Y) \geq H(X : Z). \tag{4.29}$$

Moreover, the first inequality is saturated if and only if, given Y, it is possible to reconstruct X.

This result is intuitively plausible: it tells us that if a random variable X is subject to noise, producing Y, then further actions on our part ("data processing") cannot be used to increase the amount of mutual information between the output of the process and the original information X.

Proof

The first inequality was proved in theorem 5 on page 70. From the definitions we see that $H(X : Z) \leq H(X : Y)$ is equivalent to $H(X|Y) \leq H(X|Z)$. From the fact that $X \rightarrow Y \rightarrow Z$ is a Markov chain it is easy to prove that $Z \rightarrow Y \rightarrow X$ is also a Markov chain, and thus $H(X|Y) = H(X|Y, Z)$. The problem is thus reduced to proving that $H(X, Y, Z) - H(Y, Z) = H(X|Y, Z) \leq H(X|Z) = H(X, Z) - H(Z)$. This is just the strong subadditivity inequality, which we already proved.

Suppose $H(X : Y) < H(X)$. Then it is not possible to reconstruct X from Y, since if Z is the attempted reconstruction based only on knowledge of Y, then $X \rightarrow Y \rightarrow Z$ must be a Markov process, and thus $H(X) > H(X : Z)$ by the data processing inequality. Thus $X \neq Z$. On the other hand, if $H(X : Y) = H(X)$, then we have $H(X|Y) = 0$ and thus whenever $p(X = x, Y = y) > 0$ we have $p(X = x|Y = y) = 1$. That is, if $Y = y$ then we can infer with certainty that X was equal to x, allowing us to reconstruct X.

QED

From the definition of Markov chains, it is easy to verify that if $X \rightarrow Y \rightarrow Z$ is a Markov chain, then so is $Z \rightarrow Y \rightarrow X$. Thus, as a simple corollary to the data processing inequality, we see that if $X \rightarrow Y \rightarrow Z$ is a Markov chain, then

$$H(Z : Y) \geq H(Z : X). \tag{4.30}$$

We will refer to this result as the *data pipelining inequality*. Intuitively, it says that any information Z share with X must be information which Z also shares with Y; the information is "pipelined" from X through Y to Z.

4.3 Von Neumann entropy

The Shannon entropy measures the uncertainty associated with a classical probability distribution. Quantum states are described in a similar fashion, with density operators replacing probability distributions. In this section we generalize the definition of the Shannon entropy to quantum states.

Von Neumann defined the *entropy* of a quantum state ρ by the formula

$$S(\rho) \equiv -\mathrm{tr}(\rho \log \rho). \tag{4.31}$$

In this formula logarithms are taken to base two, and we define $0 \log 0$ to be equal to zero. If λ_i are the eigenvalues of ρ then von Neumann's definition can be re-expressed

$$S(\rho) = -\sum_i \lambda_i \log \lambda_i. \tag{4.32}$$

For calculations it is usually this last formula which is most useful. For instance, the completely mixed density operator in a d dimensional space, I/d, has entropy $\log d$.

From now on, when we refer to entropy, it will usually be clear from context whether we mean Shannon or von Neumann entropy.

4.3.1 Quantum relative entropy

As for the Shannon entropy, it is extremely useful to define a quantum version of the relative entropy. Suppose ρ and σ are density operators. The *relative entropy* of ρ to σ is defined by

$$S(\rho||\sigma) \equiv \mathrm{tr}(\rho \log \rho) - \mathrm{tr}(\rho \log \sigma). \tag{4.33}$$

Conventionally, this is defined to be $+\infty$ if the kernel of σ has non-trivial intersection with the support of ρ, and is finite otherwise. The quantum relative entropy is non-negative, a result sometimes known as *Klein's inequality*.

Theorem 9 (Klein's inequality) *[182]*
The relative entropy is non-negative,

$$S(\rho||\sigma) \geq 0, \tag{4.34}$$

with equality if and only if $\rho = \sigma$.

Proof
Let $\rho = \sum_i p_i |i\rangle\langle i|$ and $\sigma = \sum_j q_j |j\rangle\langle j|$ be orthogonal decompositions for ρ and σ. From the definition of the relative entropy we have

$$S(\rho||\sigma) = \sum_i p_i \log p_i - \sum_i \langle i|\rho \log \sigma|i\rangle. \tag{4.35}$$

We substitute into this equation the equations $\langle i|\rho = p_i \langle i|$ and

$$\langle i| \log \sigma |i\rangle = \langle i| \left(\sum_j \log(q_j) |j\rangle\langle j| \right) |i\rangle = \sum_j \log(q_j) P_{ij}, \tag{4.36}$$

where $P_{ij} \equiv \langle i|j\rangle\langle j|i\rangle \geq 0$. Notice that P_{ij} satisfies the equations $\sum_i P_{ij} = 1$ and $\sum_j P_{ij} = 1$ (such matrices are called *doubly stochastic*). Substitution gives

$$S(\rho||\sigma) \quad = \quad \sum_i p_i \left(\log p_i - \sum_j P_{ij} \log(q_j) \right). \tag{4.37}$$

log is a strictly concave function, so $\sum_j P_{ij} \log q_j \leq \log r_i$, where $r_i \equiv \sum_j P_{ij} q_j$, with equality of and only if there exists a value of j for which $P_{ij} = 1$. Thus

$$S(\rho||\sigma) \quad \geq \quad \sum_i p_i \log \frac{p_i}{r_i}, \tag{4.38}$$

with equality if and only if P_{ij} is a permutation matrix. This has the form of the classical relative entropy, from which we deduce that

$$S(\rho||\sigma) \quad \geq \quad 0, \tag{4.39}$$

with equality if and only if $p_i = r_i$ for all i, and P_{ij} is a permutation matrix. To simplify the equality conditions further, note that by relabeling the eigenstates of σ if necessary, we can assume that P_{ij} is the identity matrix, and thus that ρ and σ are diagonal in the same basis. The condition $p_i = r_i$ tells us that the corresponding eigenvalues of ρ and σ are identical, and thus $\rho = \sigma$ are the equality conditions.

QED

4.3.2 Basic properties of entropy

The entropy has many interesting and useful properties.

Theorem 10 *[182]*

1. *The entropy is non-negative. The entropy is zero if and only if the state is pure.*
2. *In a d dimensional Hilbert space the entropy is at most $\log d$. The entropy is equal to $\log d$ if and only if the system is in the completely mixed state I/d.*
3. *Suppose a composite system AB is in a pure state. Then $S(A) = S(B)$.*
4. *Suppose p_i are probabilities, and ρ_i are states with mutually disjoint support. Then*

$$S(\sum_i p_i \rho_i) = H(p_i) + \sum_i p_i S(\rho_i). \tag{4.40}$$

5. **Joint entropy theorem***: Suppose p_i are probabilities, $|i\rangle$ are orthogonal states for a system A, and ρ_i is any set of density operators for another system, B. Then*

$$S(\sum_i p_i |i\rangle\langle i| \otimes \rho_i) = H(p_i) + \sum_i p_i S(\rho_i), \tag{4.41}$$

where $H(p_i)$ is the Shannon entropy of the distribution p_i.

Proof

1. Clear from the definition.
2. From the non-negativity of the relative entropy, $0 \leq S(\rho||I/d) = -S(\rho) + \log d$, from which the result follows.
3. From the Schmidt decomposition, as discussed in Appendix A, we know that the eigenvalues of systems A and system B are the same. The entropy is determined completely by the eigenvalues, so $S(A) = S(B)$.

4. Let λ_i^j and $|e_i^j\rangle$ be the eigenvalues and corresponding eigenvectors of ρ_i. Observe that $p_i\lambda_i^j$ and $|e_i^j\rangle$ are the eigenvalues and eigenvectors of $\sum_i p_i\rho_i$, and thus

$$\begin{aligned} S(\sum_i p_i\rho_i) &= -\sum_{ij} p_i\lambda_i^j \log p_i\lambda_i^j & (4.42)\\ &= -\sum_i p_i \log p_i - \sum_i p_i \sum_j \lambda_i^j \log \lambda_i^j & (4.43)\\ &= H(p_i) + \sum_i p_i S(\rho_i), & (4.44)\end{aligned}$$

   as required.

5. Immediate from the preceding result.

QED

By analogy with the Shannon entropies it is possible to define conditional and mutual von Neumann entropies. We make the definitions:

$$\begin{aligned} S(A|B) &\equiv S(A,B) - S(B) & (4.45)\\ S(A:B) &\equiv S(A) + S(B) - S(A,B) & (4.46)\\ &= S(A) - S(A|B) = S(B) - S(B|A). & (4.47)\end{aligned}$$

Some properties of the Shannon entropy fail to hold for the von Neumann entropy, and this has many interesting consequences for quantum information theory. For instance, for random variables X and Y, the inequality $H(X) \leq H(X,Y)$ holds. This makes sense: surely we cannot be more uncertain about the state of X than we are about the joint state of X and Y. This intuition fails for quantum states. Consider a system AB of two qubits in the entangled state $(|00\rangle + |11\rangle)/\sqrt{2}$. This is a pure state, so $S(A,B) = 0$. On the other hand, system A has density operator $I/2$, and thus has entropy equal to one. Another way of stating this is that for this system, the quantity $S(B|A) = S(A,B) - S(A)$ is negative.

Notice that this example involved entanglement. This is a generic feature: differences between classical and quantum information seem always to involve either or both of entanglement and the potential non-orthogonality of quantum states. For example, in Chapter 8 we will prove that the negativity of the conditional entropy always indicates that two systems are entangled, and, indeed, how negative the conditional entropy is provides a lower bound on how entangled the two systems are.

4.3.3 Measurements and entropy

How does the entropy of a quantum system behave when we perform a measurement on that system? Not surprisingly, the answer to this question depends on the type of measurement which we perform. Nevertheless, there are some surprisingly general assertions we can make about how the entropy behaves.

Suppose for example, that an orthogonal measurement described by projectors P_i is performed on a quantum system, but we never learn the result of the measurement. If the state of the system before the measurement was ρ then the state after is given by

$$\rho' \equiv \sum_i P_i \rho P_i. \quad (4.48)$$

The following result shows that the entropy is never decreased in this case, and remains the same only when the state is not changed by the measurement.

Theorem 11 *(***Orthogonal measurements increase entropy***) [182]*

Suppose P_i is a complete set of orthogonal projectors and ρ is a density operator. Then the entropy of the state $\rho' \equiv \sum_i P_i \rho P_i$ of the system after the measurement is at least as great as the original entropy,

$$S(\rho') \geq S(\rho), \tag{4.49}$$

with equality if and only if $\rho = \rho'$.

Proof *(Original?)*

The proof is to apply Klein's inequality to ρ and ρ',

$$0 \leq S(\rho'||\rho) = -S(\rho) - \mathrm{tr}(\rho \log \rho'). \tag{4.50}$$

The result will follow if we can prove that $-\mathrm{tr}(\rho \log \rho') = S(\rho')$. To do this, we apply the cyclic property of the trace and the completeness and orthogonality relations for the projectors to obtain

$$\begin{aligned} -\mathrm{tr}(\rho \log \rho') &= -\mathrm{tr}(\sum_i P_i \rho \log \rho') & (4.51) \\ &= -\mathrm{tr}(\sum_i P_i \rho \log \rho' P_i). & (4.52) \end{aligned}$$

A little thought shows that P_i commutes with ρ' and thus with $\log \rho'$, so

$$\begin{aligned} -\mathrm{tr}(\rho \log \rho') &= -\mathrm{tr}(\sum_i P_i \rho P_i \log \rho') & (4.53) \\ &= -tr(\rho' \log \rho') = S(\rho'). & (4.54) \end{aligned}$$

This completes the proof.

QED

4.3.4 The entropy of ensembles

Theorem 12 *[182]*

Suppose $\rho = \sum_i p_i \rho_i$, where p_i are some set of probabilities, and the ρ_i are density operators. Then

$$S(\rho) \leq H(p_i) + \sum_i p_i S(\rho_i), \tag{4.55}$$

with equality if and only if the states ρ_i have support on orthogonal subspaces.

Proof

We begin with the pure state case, $\rho_i = |\psi_i\rangle\langle\psi_i|$. Let A be a system with the same state space as the ρ_i, and introduce a system B with an orthonormal basis $|i\rangle$ corresponding to the index i on the probabilities p_i. Define

$$|AB\rangle \equiv \sum_i \sqrt{p_i} |\psi_i\rangle |i\rangle. \tag{4.56}$$

Since $|AB\rangle$ is a pure state we have

$$S(B) = S(A) = S(\sum_i p_i |\psi_i\rangle\langle\psi_i|) = S(\rho). \tag{4.57}$$

Suppose we perform an orthonormal measurement on the system B in the $|i\rangle$ basis. After the measurement the state of system B is

$$B' = \sum_i p_i |i\rangle\langle i|. \tag{4.58}$$

But orthogonal measurements never decrease entropy, so $S(\rho) = S(B) \leq S(B') = H(p_i)$. Observing that $S(\rho_i) = 0$ for the pure state case, we have proved that

$$S(\rho) \leq H(p_i) + \sum_i p_i S(\rho_i), \tag{4.59}$$

when the states ρ_i are pure states. Furthermore, equality holds if and only if $B = B'$, which is easily seen to occur if and only if the states $|\psi_i\rangle$ are orthogonal.

The mixed state case is now easy. Let $\rho_i = \sum_j p_j^i |e_j^i\rangle\langle e_j^i|$ be orthonormal decompositions for the states ρ_i, so $\rho = \sum_{ij} p_i p_j^i |e_j^i\rangle\langle e_j^i|$. Applying the pure state result and the observation that $\sum_j p_j^i = 1$ for each i, we have

$$\begin{aligned} S(\rho) &\leq -\sum_{ij} p_i p_j^i \log(p_i p_j^i) && (4.60)\\ &= -\sum_i p_i \log p_i - \sum_i p_i \sum_j p_j^i \log p_j^i && (4.61)\\ &= H(p_i) + \sum_i p_i S(\rho)_i, && (4.62) \end{aligned}$$

which is the desired result. The equality conditions for the mixed state case follow immediately from the equality conditions for the pure state case.

QED

4.3.5 Subadditivity

Suppose distinct quantum systems, A and B, have a joint state ρ_{AB}. Then the joint entropy for the two systems satisfies the inequalities

$$\begin{aligned} S(A,B) &\leq S(A) + S(B) && (4.63)\\ S(A,B) &\geq |S(A) - S(B)|. && (4.64) \end{aligned}$$

The first of these inequalities is known as the *subadditivity* inequality for Von Neumann entropy, and holds with equality if and only if systems A and B are uncorrelated, that is, $\rho_{AB} = \rho_A \otimes \rho_B$. The second is called the *triangle* inequality, or sometimes the *Araki-Lieb* inequality.

The proof of subadditivity is a simple application of Klein's inequality, $S(\rho) \leq -\mathrm{tr}(\rho \log \sigma)$. Setting $\rho \equiv \rho_{AB}$ and $\sigma \equiv \rho_A \otimes \rho_B$, note that

$$\begin{aligned} -\mathrm{tr}(\rho \log \sigma) &= -\mathrm{tr}(\rho_{AB}(\log \rho_A + \log \rho_B)) && (4.65)\\ &= -\mathrm{tr}(\rho_A \log \rho_A) - \mathrm{tr}(\rho_B \log \rho_B) && (4.66)\\ &= S(\rho_A) + S(\rho_B). && (4.67) \end{aligned}$$

Klein's inequality therefore gives $S(\rho_{AB}) \leq S(\rho_A) + S(\rho_B)$, as desired. The equality conditions $\rho = \sigma$ for Klein's inequality give equality conditions $\rho_{AB} = \rho_A \otimes \rho_B$ for subadditivity.

To prove the triangle inequality, let R be a system which purifies systems A and B[1]. Applying subadditivity we have

$$S(R) + S(A) \geq S(A, R). \tag{4.68}$$

Since ABR is in a pure state, $S(A, R) = S(B)$ and $S(R) = S(A, B)$. The previous inequality then may be rearranged to give

$$S(A, B) \geq S(B) - S(A). \tag{4.69}$$

The equality conditions for this inequality are not so easy to understand. Formally, the equality conditions are that $\rho_{AR} = \rho_A \otimes \rho_R$. Intuitively, what this means is that A is already as entangled as it can possibly be with the outside world, given its existing correlations with system B. Note also that by symmetry between the systems A and B we also have, $S(A, B) \geq S(A) - S(B)$. Combining these two inequalities gives the triangle inequality.

4.3.6 Concavity of the entropy

The entropy is a *concave* function of its inputs. That is, given real numbers λ_i satisfying $\lambda_i \geq 0, \sum_i \lambda_i = 1$, and corresponding density operators ρ_i, the entropy satisfies the equation:

$$S(\sum_i \lambda_i \rho_i) \geq \sum_i \lambda_i S(\rho_i). \tag{4.70}$$

To understand why this should be so, imagine that the λ_is are probabilities. Then $\sum_i \lambda_i \rho_i$ expresses the state of a quantum system which is in an unknown state ρ_i with probability λ_i. Not surprisingly, our uncertainty about this mixture of states should be higher than the average uncertainty of the states ρ_i.

Let A have a state space containing the state ρ_i, and let B have a state space with orthonormal basis $|i\rangle$. Define the joint state

$$\rho^{AB} \equiv \sum_i \lambda_i \rho_i \otimes |i\rangle\langle i|. \tag{4.71}$$

To prove concavity we use the subadditivity of the entropy. Note that

$$S(A) = S(\sum_i \lambda_i \rho_i) \tag{4.72}$$

$$S(B) = S(\sum_i \lambda_i |i\rangle\langle i|) = H(\lambda_i) \tag{4.73}$$

$$S(AB) = H(\lambda_i) + \sum_i \lambda_i S(\rho_i). \tag{4.74}$$

Applying the inequality $S(AB) \leq S(A) + S(B)$ we obtain

$$\sum_i \lambda_i S(\rho_i) \leq S(\sum_i \lambda_i \rho_i), \tag{4.75}$$

[1]See Appendix A for a review of purifications. R purifies A and B if the joint state of RAB is pure.

which is the desired concavity result. Note that equality holds if and only if all the states ρ_i are identical; that is, the entropy is a strictly concave function[2] of its inputs.

It's worth pausing here to think about the strategy we've employed in this proof, and the similar strategy used to prove the triangle inequality. We introduced an *auxiliary system*, B, in order to prove a result about the system A. Introducing auxiliary systems is something often done in quantum information theory, and we'll see this trick again and again. The intuition behind the introduction of B in this particular manner is as follows: we want to find a system part of which is in the state $\sum_i \lambda_i \rho_i$, where the value of i is not known. System B effectively stores the "true" value of i: if A were "truly" in state ρ_i, the system B would be in state $|i\rangle\langle i|$, and observing system B in the $|i\rangle$ basis would reveal this fact. Using auxiliary systems in this way to encode our intuition in a rigorous way is something of an art, but it is also an essential part of many proofs in quantum information theory.

4.4 Strong subadditivity

The subadditivity inequalities proved in the last section for two quantum systems can be extended to three systems. The basic result is known as the *strong subadditivity* inequality, and it is one of the most important and useful results in quantum information theory. Unfortunately, unlike in the classical case, proving the quantum strong subadditivity inequality appears to be quite difficult. However, it will be used frequently throughout this Dissertation, so we give a full proof here. The result was first proved by Lieb and Ruskai [115], based upon an earlier result of Lieb [114]. The proof of Lieb's theorem given here is adapted from Bhatia [26], which is an adaptation of a proof of Simon [168].

The *strong subadditivity inequality* for von Neumann entropies states that for a trio of quantum systems, A, B, C,

$$S(A,B,C) + S(B) \leq S(A,B) + S(B,C). \tag{4.76}$$

The proof of this inequality which we give is based upon a deep mathematical result known as *Lieb's theorem*. We begin with a few simple notations and definitions.

Suppose $f(A,B)$ is a real valued function of two matrices, A and B. Then f is said to be *jointly concave* in A and B if for all $0 \leq \lambda \leq 1$,

$$f(\lambda A_1 + (1-\lambda)A_2, \lambda B_1 + (1-\lambda)B_2) \geq \lambda f(A_1, B_1) + (1-\lambda) f(A_2, B_2). \tag{4.77}$$

For matrices A and B, we say $A \geq B$ if $A - B$ is a positive matrix. If A is a positive matrix, and t a real number, then we define A^t as follows. Let $A = UDU^\dagger$, where U is unitary and D is a diagonal matrix with non-negative entries. Define D^t to be the diagonal matrix with entries d_i^t, where d_i are the diagonal entries in D. Define $A^t \equiv UD^tU^\dagger$. Let A be an arbitrary matrix. We define the *norm* of A by

$$\|A\| \equiv \max_{\langle u|u\rangle = 1} |\langle u|A|u\rangle|. \tag{4.78}$$

In our proof of Lieb's theorem and strong subadditivity, we will have occasion to use the following easily verified observations.

[2]This observation can be used to give an elegant proof that the unique maximal entropy state is the completely mixed state. Let ρ be given, and note that $I/d = \sum_\pi \rho_\pi / d!$, where the sum is over all permutations π on d elements, and ρ_π is obtained from ρ by a permutation of the basis elements in which ρ is diagonal. The result follows by strict concavity.

1. If $A \leq B$, then $XAX^\dagger \leq XBX^\dagger$ for all matrices X.
2. Let $f(A, B)$ be a jointly concave function. Then $f(A, B)$ is concave in A, with B held fixed. It is easy to find a function of two variables that is concave in each of its inputs, but is not jointly concave.
3. $A \geq 0$ if and only if A is a positive operator.
4. The relation $\geq$ is a partial order on operators – that is, it is transitive ($A \geq B$ and $B \geq C$ implies $A \geq C$), asymmetric ($A \geq B$ and $B \geq A$ implies $A = B$), and reflexive ($A \geq A$).
5. Suppose A has eigenvalues λ_i. Define λ to be the maximum of the set $|\lambda_i|$. Then:
   (a) $\|A\| \geq \lambda$.
   (b) When A is Hermitian, $\|A\| = \lambda$.
   (c) When
   $$A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \tag{4.79}$$
   it is easy to verify the $\|A\| = 3/2 > 1 = \lambda$.
6. The eigenvalues of A are the solutions to the characteristic equation $\det(xI - A) = 0$. For invertible A, note that $\det(xI - AB) = \det A \det(xI - BA) \det A^{-1} = \det(xI - BA)$, and thus the eigenvalues of AB and BA are the same. A simple continuity argument shows that this is generally true in finite dimensions.
7. Suppose A and B are such that AB is Hermitian. Then from the previous two observations it follows that $\|AB\| \leq \|BA\|$.
8. Suppose A is positive. Then $\|A\| \leq 1$ if and only if $A \leq I$.
9. Let A be a positive matrix. Define a superoperator (linear operator on matrices) by the equation $\mathcal{A}(X) \equiv AX$. Then $\mathcal{A}$ is positive with respect to the Hilbert-Schmidt inner product. That is, for all X, $\mathrm{tr}(X^\dagger \mathcal{A}(X)) \geq 0$. Similarly, the superoperator defined by $\mathcal{A}(X) \equiv XA$ is positive with respect to the Hilbert-Schmidt inner product on matrices.

With these results in hand, we are now in a position to state and prove Lieb's theorem.

Theorem 13 (Lieb's theorem) *[114]*
Let X be a matrix, and $0 \leq t \leq 1$. Then the function

$$f(A, B) \equiv tr(X^\dagger A^t X B^{1-t}) \tag{4.80}$$

is jointly concave in positive matrices A and B.

Lieb's theorem is an easy corollary of the following lemma:

Lemma 1 *Let $R_1, R_2, S_1, S_2, T_1, T_2$ be positive operators such that $0 = [R_1, R_2] = [S_1, S_2] = [T_1, T_2]$, and*

$$R_1 \geq S_1 + T_1 \tag{4.81}$$
$$R_2 \geq S_2 + T_2. \tag{4.82}$$

Then for all $0 \leq t \leq 1$*,*

$$R_1^t R_2^{1-t} \geq S_1^t S_2^{1-t} + T_1^t T_2^{1-t} \tag{4.83}$$

is true as a matrix inequality.

Proof (*Adapted from [26]*)

We begin by proving the result for $t = 1/2$, and then use this to establish the result for general t.

Let $|x\rangle$ and $|y\rangle$ be any two vectors. Applying the Cauchy-Schwartz inequality twice and some straightforward manipulations, we have

$$\begin{aligned}
& |\langle x|(S_1^{1/2}S_2^{1/2} + T_1^{1/2}T_2^{1/2})|y\rangle| \\
\leq\ & |\langle x|S_1^{1/2}S_2^{1/2}|y\rangle| + |\langle x|T_1^{1/2}T_2^{1/2}|y\rangle| && (4.84)\\
\leq\ & \|S_1^{1/2}|x\rangle\|\,\|S_2^{1/2}|y\rangle\| + \|T_1^{1/2}|x\rangle\|\,\|T_2^{1/2}|y\rangle\| && (4.85)\\
\leq\ & \sqrt{\left(\|S_1^{1/2}|x\rangle\|^2 + \|T_1^{1/2}|x\rangle\|^2\right)\left(\|S_2^{1/2}|y\rangle\|^2 + \|T_2^{1/2}|y\rangle\|^2\right)} && (4.86)\\
=\ & \sqrt{\langle x|(S_1+T_1)|x\rangle\langle y|(S_2+T_2)|y\rangle}. && (4.87)
\end{aligned}$$

By hypothesis, $S_1 + T_1 \leq R_1$ and $S_2 + T_2 \leq R_2$, so

$$|\langle x|(S_1^{1/2}S_2^{1/2} + T_1^{1/2}T_2^{1/2})|y\rangle| \leq \sqrt{\langle x|R_1|x\rangle\langle y|R_2|y\rangle}. \tag{4.88}$$

Let $|u\rangle$ be any unit vector. Then applying the previous result with $|x\rangle \equiv R_1^{-1/2}|u\rangle$ and $|y\rangle \equiv R_2^{-1/2}|u\rangle$ gives

$$\begin{aligned}
& \langle u|R_1^{-1/2}(S_1^{1/2}S_2^{1/2} + T_1^{1/2}T_2^{1/2})R_2^{-1/2}|u\rangle \\
\leq\ & \sqrt{\langle u|R_1^{-1/2}R_1R_1^{-1/2}|u\rangle\langle u|R_2^{-1/2}R_2R_2^{-1/2}|u\rangle} && (4.89)\\
=\ & \sqrt{\langle u|u\rangle\langle u|u\rangle} = 1. && (4.90)
\end{aligned}$$

Thus

$$\|R_1^{-1/2}(S_1^{1/2}S_2^{1/2} + T_1^{1/2}T_2^{1/2})R_2^{-1/2}\| \leq 1. \tag{4.91}$$

Define

$$\begin{aligned}
A &\equiv R_1^{-1/4}R_2^{-1/4}(S_1^{1/2}S_2^{1/2} + T_1^{1/2}T_2^{1/2})R_2^{-1/2} && (4.92)\\
B &\equiv R_2^{1/4}R_1^{-1/4}. && (4.93)
\end{aligned}$$

Note that AB is Hermitian, so by observation number 7 on page 81,

$$\begin{aligned}
& \|R_1^{-1/4}R_2^{-1/4}(S_1^{1/2}S_2^{1/2} + T_1^{1/2}T_2^{1/2})R_2^{-1/4}R_1^{-1/4}\| \\
=\ & \|AB\| \leq \|BA\| && (4.94)\\
=\ & \|R_1^{-1/2}(S_1^{1/2}S_2^{1/2} + T_1^{1/2}T_2^{1/2})R_2^{-1/2}\| && (4.95)\\
\leq\ & 1, && (4.96)
\end{aligned}$$

where the last inequality is just (4.91). AB is a positive operator, so by observation number 8 on page 81 and the previous inequality,

$$R_1^{-1/4}R_2^{-1/4}(S_1^{1/2}S_2^{1/2}+T_1^{1/2}T_2^{1/2})R_2^{-1/4}R_1^{-1/4}\leq I. \tag{4.97}$$

Finally, by observation 1 on page 81, and the commutativity of R_1 and R_2,

$$S_1^{1/2}S_2^{1/2}+T_1^{1/2}T_2^{1/2}\leq R_1^{1/2}R_2^{1/2}, \tag{4.98}$$

which establishes that (4.83) holds for $t=1/2$.

Let I be the set of all t such that (4.83) holds. By inspection, we see that 0 and 1 are elements of I. We now use the result for $t=1/2$ to prove the result for general t. Suppose μ and η are elements of I, so

$$R_1^{\mu}R_2^{1-\mu} \geq S_1^{\mu}S_2^{1-\mu}+T_1^{\mu}T_2^{1-\mu} \tag{4.99}$$
$$R_1^{\eta}R_2^{1-\eta} \geq S_1^{\eta}S_2^{1-\eta}+T_1^{\eta}T_2^{1-\eta}. \tag{4.100}$$

These inequalities are of the form (4.81) and (4.82) for which the $t=1/2$ case has already been proved. Using the $t=1/2$ result we see that

$$\left(R_1^{\mu}R_2^{1-\mu}\right)^{1/2}\left(R_1^{\eta}R_2^{1-\eta}\right)^{1/2}\geq\left(S_1^{\mu}S_2^{1-\mu}\right)^{1/2}\left(S_1^{\eta}S_2^{1-\eta}\right)^{1/2}+\left(T_1^{\mu}T_2^{1-\mu}\right)^{1/2}\left(T_1^{\eta}T_2^{1-\eta}\right)^{1/2}. \tag{4.101}$$

Using the commutativity assumptions $0=[R_1,R_2]=[S_1,S_2]=[T_1,T_2]$, we see that for $\nu\equiv(\mu+\eta)/2$,

$$R_1^{\nu}R_2^{1-\nu}\geq S_1^{\nu}S_2^{1-\nu}+T_1^{\nu}T_2^{1-\nu}. \tag{4.102}$$

Thus whenever μ and η are in I, so is $(\mu+\eta)/2$. Since 0 and 1 are in I, it is easy to see that any number x between 0 and 1 with a finite binary expansion must be in I. Thus I is dense in $[0,1]$. The result now follows from the continuity in t of the conclusion, (4.83).

QED

The proof of Lieb's theorem is a simple application of the lemma. The main novelty is that the operators in the lemma are chosen to be *superoperators* – linear maps on operators. These will be chosen in such a way as to be positive with respect to the Hilbert-Schmidt inner product $(A,B)\equiv\mathrm{tr}(A^{\dagger}B)$.

Proof (Lieb's theorem) *(Adapted from [26])*

Define

$$\mathcal{S}_1(X) \equiv \lambda A_1 X \tag{4.103}$$
$$\mathcal{S}_2(X) \equiv \lambda X B_1 \tag{4.104}$$
$$\mathcal{T}_1(X) \equiv (1-\lambda)A_2 X \tag{4.105}$$
$$\mathcal{T}_2(X) \equiv (1-\lambda)X B_2 \tag{4.106}$$
$$\mathcal{R}_1 \equiv \mathcal{S}_1+\mathcal{T}_1 \tag{4.107}$$
$$\mathcal{R}_2 \equiv \mathcal{S}_2+\mathcal{T}_2. \tag{4.108}$$

Observe that $\mathcal{S}_1$ and $\mathcal{S}_2$ commute, as do $\mathcal{T}_1$ and $\mathcal{T}_2$, and $\mathcal{R}_1$ and $\mathcal{R}_2$. Recall observation 9 on page 81, that all these operators are positive with respect to the Hilbert-Schmidt inner product. By the lemma,

$$\mathcal{R}_1^t\mathcal{R}_2^{1-t}\geq\mathcal{S}_1^t\mathcal{S}_2^{1-t}+\mathcal{T}_1^t\mathcal{T}_2^{1-t}. \tag{4.109}$$

Taking the $X \cdot X$ matrix element of the previous inequality gives

$$\begin{aligned} & \mathrm{tr}\left[X^\dagger \left(\lambda A_1 + (1-\lambda)A_2\right)^t X \left(\lambda B_1 + (1-\lambda)B_2\right)^{1-t}\right] \\ \geq\ & \mathrm{tr}\left[X^\dagger (\lambda A_1)^t X (\lambda B_1)^{1-t}\right] + \mathrm{tr}\left[X^\dagger ((1-\lambda)A_2)^t X ((1-\lambda)B_2)^{1-t}\right] \qquad (4.110) \\ =\ & \lambda \mathrm{tr}(X^\dagger A_1^t X B_1^{1-t}) + (1-\lambda)\mathrm{tr}(X^\dagger A_2^t X B_2^{1-t}), \qquad (4.111) \end{aligned}$$

which is the desired statement of joint concavity.

QED

Let B and C be density operators. Recall that the *relative entropy* of B to C is defined by

$$S(B\|C) \equiv -S(B) - \mathrm{tr}(B \log C). \qquad (4.112)$$

Theorem 14 (Convexity of the relative entropy) *[182]*
The relative entropy $S(B\|C)$ is jointly convex in its arguments.

Proof [26]
Define

$$I_t(A, X) \equiv \mathrm{tr}(X^\dagger A^t X A^{1-t}) - \mathrm{tr}(X^\dagger X A). \qquad (4.113)$$

Note that the first term in this expression is concave in A, by Lieb's theorem, and the second term is linear in A. Thus, $I_t(A, X)$ is concave in A. Define

$$I(A, X) \equiv \frac{d}{dt}\,|_{t=0}\ I_t(A, X) = \mathrm{tr}(X^\dagger (\log A) X A) - \mathrm{tr}(X^\dagger X (\log A) A). \qquad (4.114)$$

Noting that $I_0(A, X) = 0$ and using the concavity of $I_t(A, X)$ in A we have

$$\begin{aligned} I(\lambda A_1 + (1-\lambda)A_2, X) &= \lim_{\Delta \to 0} \frac{I_\Delta(\lambda A_1 + (1-\lambda)A_2, X)}{\Delta} \qquad (4.115) \\ &\geq \lambda \lim_{\Delta \to 0} \frac{I_\Delta(A_1, X)}{\Delta} + (1-\lambda) \lim_{\Delta \to 0} \frac{I_\Delta(A_2, X)}{\Delta} \qquad (4.116) \\ &= \lambda I(A_1, X) + (1-\lambda) I(A_2, X). \qquad (4.117) \end{aligned}$$

That is, $I(A, X)$ is a concave function of A.

Finally, defining the block matrices

$$A \equiv \begin{bmatrix} B & 0 \\ 0 & C \end{bmatrix}, \quad X \equiv \begin{bmatrix} 0 & 0 \\ I & 0 \end{bmatrix} \qquad (4.118)$$

we can easily verify that $I(A, X) = -S(B\|C)$. The joint convexity of $S(B\|C)$ now follows from the concavity of $I(A, X)$ in A.

QED

Suppose ρ_{AB} is the state of a joint system, AB. Recall the definition of the *conditional entropy* of system A given system B,

$$S(A|B) \equiv S(A, B) - S(B). \qquad (4.119)$$

Corollary 1 *$S(A|B)$ is concave in ρ_{AB}.*

Proof
Let d be the dimension of system A. Note that

$$\begin{aligned} S(\rho_{AB}\|\frac{I}{d}\otimes\rho_B) &= -S(A,B) - \mathrm{tr}(\rho_{AB}\log(\frac{I}{d}\otimes\rho_B)) && (4.120)\\ &= -S(A,B) - \mathrm{tr}(\rho_B\log\rho_B) + \log d && (4.121)\\ &= -S(A|B) + \log d. && (4.122)\end{aligned}$$

Thus $S(A|B) = \log d - S(\rho_{AB}\|I/d\otimes\rho_B)$. The concavity of $S(A|B)$ follows from the joint convexity of the relative entropy.
QED
Strong subadditivity can now be proved using the convexity of the conditional entropy.

Theorem 15 *[115]*
For any trio of quantum systems, A, B, C, the inequalities

$$\begin{aligned} S(A)+S(B) &\leq S(A,C)+S(B,C) && (4.123)\\ S(A,B,C)+S(B) &\leq S(A,B)+S(B,C) && (4.124)\end{aligned}$$

hold.

Proof [115]
The two inequalities are, in fact, equivalent. We will use convexity of the conditional entropy to prove the first, and show that the second follows. Define a function of density operators on the system ABC,

$$T(\rho_{ABC}) \equiv S(A)+S(B)-S(A,C)-S(B,C) = -S(C|A)-S(C|B). \quad (4.125)$$

From the concavity of the relative entropy we see that $T(\rho_{ABC})$ is a convex function of ρ_{ABC}. Let $\rho_{ABC} = \sum_i p_i P_i$, where P_i is a pure state of the system ABC and the p_i are probabilities. From the convexity of T, $T(\rho_{ABC}) \leq \sum_i p_i T(P_i)$. But for a pure state, $T(P_i) = 0$, as $S(A,C) = S(B)$ and $S(B,C) = S(A)$ for a pure state. It follows that $T(\rho_{ABC}) \leq 0$, and thus

$$S(A)+S(B)-S(A,C)-S(B,C) \leq 0, \quad (4.126)$$

which is the first inequality.
Finally, to obtain the second inequality, introduce a fourth system, R, purifying the system ABC. Then

$$S(R)+S(B) \leq S(R,C)+S(B,C). \quad (4.127)$$

Since $ABCR$ is a pure state, $S(R) = S(A,B,C)$ and $S(R,C) = S(A,B)$, so the previous inequality becomes

$$S(A,B,C)+S(B) \leq S(A,B)+S(B,C), \quad (4.128)$$

as we set out to show.
QED
Strong subadditivity and the convexity of the relative entropy are results which have many useful consequences. We will use these results many, many times throughout the remainder of this Dissertation. For the time being, it is interesting to note a few elementary consequences.

First, it is worth emphasizing how remarkable it is that the inequality $S(A) + S(B) \leq S(A,C) + S(B,C)$ holds. The corresponding inequality holds also for Shannon entropies, but for quite different reasons. For Shannon entropies it is true that $H(A) \leq H(A,C)$ and $H(B) \leq H(B,C)$, so the sum of the two inequalities must necessarily be true. In the quantum case, it is possible to have either $S(A) > S(A,C)$ or $S(B) > S(B,C)$, yet somehow nature manages to conspire in such a way that both these possibilities are not true simultaneously, in order to ensure that the condition $S(A) + S(B) \leq S(A,C) + S(B,C)$ is always satisfied. Other ways of rephrasing this are in terms of conditional entropies and mutual informations,

$$0 \leq S(C|A) + S(C|B) \tag{4.129}$$

$$S(A:B) + S(A:C) \leq 2S(A), \tag{4.130}$$

both of which are also remarkable inequalities, for similar reasons. Note, however, that the inequality $0 \leq S(A|C) + S(B|C)$, which one might hope to be true based upon (4.129) is not, as can easily be seen by choosing BC to be a Bell state in a product state with system A. In part, it is these wonderful facts which brought home to me how strange and counter-intuitive quantum entropies may be.

There is an interesting set of questions related to the subadditivity properties of quantum conditional entropies. We already saw earlier that the Shannon mutual information is not subadditive, and thus the quantum mutual information is not subadditive, either. What about the subadditivity of the conditional entropy? That is, is it true that

$$S(A_1, A_2|B_1, B_2) \leq S(A_1|B_1) + S(A_2|B_2), \tag{4.131}$$

for any four quantum systems A_1, A_2, B_1 and B_2? It turns out that this inequality is correct. To prove this, we apply the strong subadditivity inequality.

By strong subadditivity,

$$S(A_1, A_2, B_1, B_2) + S(B_1) \leq S(A_1, B_1) + S(A_2, B_1, B_2). \tag{4.132}$$

Adding $S(B_2)$ to each side of this inequality, we obtain

$$S(A_1, A_2, B_1, B_2) + S(B_1) + S(B_2) \leq S(A_1, B_1) + S(A_2, B_1, B_2) + S(B_2). \tag{4.133}$$

Applying strong subadditivity to the last two terms of the right hand side gives

$$S(A_1, A_2, B_1, B_2) + S(B_1) + S(B_2) \leq S(A_1, B_1) + S(A_2, B_2) + S(B_1, B_2). \tag{4.134}$$

Rearranging this inequality gives

$$S(A_1, A_2|B_1, B_2) \leq S(A_1|B_1) + S(A_2|B_2), \tag{4.135}$$

which is the desired statement of subadditivity of the conditional entropy.

Two closely related results are the subadditivity of the conditional entropy in the first and second entries. These results are attributed by Ruskai [149] to work of Lieb, which I have not been able to locate. For example, subadditivity in the first entry, $S(A, B|C) \leq S(A|C) + S(B|C)$ is trivially seen to be equivalent to strong subadditivity. Subadditivity in the second entry is slightly more difficult to prove. We wish to show that $S(A|B, C) \leq S(A|B) + S(A|C)$. Note that this is equivalent to demonstrating the inequality

$$S(A, B, C) + S(B) + S(C) \leq S(A, B) + S(B, C) + S(A, C). \tag{4.136}$$

To prove this note that at least one of the inequalities $S(C) \leq S(A,C)$ or $S(B) \leq S(A,B)$ must be true, as $S(A|B) + S(A|C) \geq 0$. Suppose $S(C) \leq S(A,C)$. Adding to this inequality the strong subadditivity inequality, $S(A,B,C) + S(B) \leq S(A,B) + S(B,C)$ gives the result. A similar proof holds in the case when $S(B) \leq S(A,B)$.

We will return to the subadditivity properties of conditional information in Chapter 10, in the context of the quantum channel capacity, where they play a crucial role in understanding what is going on.

To finish the Chapter, let us look a little more closely at the convexity of the relative entropy. Earlier we defined the relative entropy for density operators, however there is no reason we cannot extend the definition to any two positive operators, A and B,

$$S(A||B) \equiv -S(A) + \mathrm{tr}(A \log B), \tag{4.137}$$

with the same conventions as before. Following the earlier argument used to establish the convexity of the relative entropy we see that the general relative entropy is also convex. This has an interesting consequence for the case of density operators, although I am yet to find any practical use for the pretty theorem we will shortly prove.

Ruskai [148] has pointed out the following interesting homogeneity relation,

$$S(\alpha A||\beta B) = \alpha S(A||B) + \alpha \mathrm{tr}(A) \log(\alpha/\beta), \tag{4.138}$$

which holds for $\alpha, \beta > 0$. Note that when $\alpha = \beta$ we deduce that $S(\alpha A||\alpha B) = \alpha S(A||B)$.

This is an observation with many interesting consequences. First, we see that to prove the double convexity of the relative entropy, it suffices to prove that

$$S(A_1 + A_2||B_1 + B_2) \leq S(A_1||B_1) + S(A_2||B_2). \tag{4.139}$$

We might refer to this inequality as the "joint subadditivity" of the relative entropy; we will see below that it also follows from joint convexity, so the two statements are equivalent. If joint subadditivity were true, then we would have

$$\begin{aligned} S(\lambda A_1 + (1-\lambda)A_2||\lambda B_1 + (1-\lambda)B_2) &\leq S(\lambda A_1||\lambda B_1) + S((1-\lambda)A_2||(1-\lambda)B_2) && (4.140) \\ &= \lambda S(A_1||B_1) + (1-\lambda)S(A_2||B_2), && (4.141) \end{aligned}$$

which is the desired double convexity result. The reason I mention this is with a viewpoint to future proofs of strong subadditivity: it may be that it is easier to try proving the joint subadditivity property of the relative entropy, rather than attempting the joint convexity directly.

Next we turn this result around and see that the joint subadditivity follows from joint convexity, that is,

$$S(\sum_i A_i||\sum_i B_i) \leq \sum_i S(A_i||B_i) \tag{4.142}$$

is itself a consequence of the double convexity, since if i ranges over n indices, then

$$\begin{aligned} S(\sum_i A_i||\sum_i B_i) &= S(\frac{1}{n}\sum_i (nA_i)||\frac{1}{n}\sum_i (nB_i)) && (4.143) \\ &\leq \sum_i \frac{S(nA_i||nB_i)}{n} && (4.144) \\ &= \sum_i S(A_i||B_i). && (4.145) \end{aligned}$$

This circle of ideas can be combined to give a new and rather pretty convexity result for relative entropy. The proof is immediate from Ruskai's homogeneity relation and the convexity of the relative entropy:

Theorem 16 *Let p_i and q_i be probability distributions over the same set of indices. Then*

$$S(\sum_i p_i A_i || \sum_i q_i B_i) \leq \sum_i p_i S(A_i||B_i) + \sum_i p_i tr(A_i) \log(p_i/q_i). \tag{4.146}$$

In the case where the A_i are density operators so $tr(A_i) = 1$, this reduces to the remarkable formula

$$S(\sum_i p_i A_i || \sum_i q_i B_i) \leq \sum_i p_i S(A_i||B_i) + H(p_i||q_i), \tag{4.147}$$

where $H(\cdot||\cdot)$ is the Shannon relative entropy.

Summary of Chapter 4: Entropy and information

- **Fundamental measures of information arise as the answers to fundamental questions about the quantity of physical resources required to solve some information processing problem.**
- **Basic definitions:**

$$\begin{aligned} S(A) &\equiv -\mathrm{tr}(A \log A) \quad \textit{(entropy)} && (4.148)\\ S(A||B) &\equiv -S(A) + \mathrm{tr}(A \log B) \quad \textit{(relative entropy)} && (4.149)\\ S(A|B) &\equiv S(A,B) - S(B) \quad \textit{(conditional entropy)} && (4.150)\\ S(A:B) &\equiv S(A) + S(B) - S(A,B) \quad \textit{(mutual information).} && (4.151) \end{aligned}$$

- **The relative entropy is jointly convex in its arguments.**
- **Strong subadditivity:** $S(A,B,C) + S(B) \leq S(A,B) + S(B,C)$. The other entropy inequalities we discussed are corollaries of this or the joint convexity of the relative entropy.

Chapter 5

Distance measures for quantum information

What does it mean to say that information is preserved during some process? What does it mean to say that two items of information are similar? A well developed theory of quantum information must provide useful answers to these questions. Because of the wide variety of information types in quantum mechanics, inequivalent answers to these questions are possible, with each answer useful in the context of a specific class of information processing tasks.

The principle concern of this Chapter is the development of *distance measures* for quantum information. We will be concerned with two classes of distance measures, *static measures*, and *dynamic measures*. Static measures provide a quantitative means of determining how close two quantum states are, while dynamic measures provide a quantitative means of determining how well information has been preserved during a dynamic process. The strategy used in this Chapter is to begin by developing good static measures of distance, and then to use those static measures to aid in the development of good dynamic measures of distance.

Each of the distance measures introduced in this Chapter can be viewed in two ways. First, and most important, we inquire as to the *operational meaning* of the distance measures. That is, we attempt to find a physical question which leads naturally to that distance measure. For example, one of the measures we introduce, the absolute distance, turns out to be directly related to the ability to distinguish two quantum states by measurements on those states. Second, distance measures can be viewed as purely mathematical constructs, useful for proving facts about the behaviour of quantum systems. For example, we will introduce a quantity, the fidelity, which does not appear to have an especially clear physical meaning. However, properties of the fidelity can be used to prove facts of great physical significance, such as the existence of unique stationary states for certain open quantum systems.

The primary purpose of the Chapter is to serve as review and reference for basic properties of the distance measures we will consider, however, it is also contains numerous results which I am not aware of elsewhere in the literature. In places the Chapter contains rather detailed mathematics; upon a first read, these sections may be read lightly, and returned to later for reference purposes.

5.1 Distance measures for classical information

We begin by studying distance measures for classical information. In the classical setting, we will discuss three primary notions of distance. Two of these notions will be *static measures* of distance, involving the comparison of two classical probability distributions. The third notion of distance which we examine is a *dynamical measure* of distance, which is associated with a process. We won't prove many general results in this section, because we will prove quantum generalizations of the classical results later in the Chapter. Indeed, the discussion in this section may appear rather trivial, and the reader might wonder why we didn't simply skip to the quantum case. However, the intuitive justification for the distance measures is easier to grasp in the classical situation, justifying a separate treatment.

What are the objects to be compared in classical information theory? In some circumstances it is useful to compare strings of bits. For that purpose, the *Hamming distance* is perhaps the most commonly used measure of distance; it is defined to be the number of places at which two bit strings are not equal. In this Dissertation we will have little concern with the actual labeling of bit strings, so notions such as Hamming distance are of little interest to us.

By contrast, we will be very concerned with the comparison of information sources. In classical information theory, an information source is usually modeled as a random variable, or equivalently, a probability distribution, over some source alphabet. For example, an unknown source of English text may be modeled as a sequence of random variables over the Roman alphabet. Before the text is read, we can make a fair guess at the relative frequency of the letters that will appear in the text, and certain correlations among them, such as fact that occurrences of the pair of letters "th" are much more common than the pair "zx" in English text. This characterization of information sources as probability distributions over some alphabet causes us to concentrate on the comparison of probability distributions in our search for measures of distance.

What does it mean to say that two classical probability distributions, p_x and q_x, over the same index set, x, are near to one another? It is difficult to give an answer to this question which is obviously the unique "correct" answer, so instead we will propose several different answers, each of which is useful in particular contexts.

The first measure is the L_1 distance, defined by the equation

$$D(p_x, q_x) \equiv \sum_x |p_x - q_x|. \tag{5.1}$$

More usually, we will refer to this as the *absolute distance* between the probability distributions p_x and q_x. The absolute distance is easily seen to be a metric on probability distributions, so the use of the term "distance" is justified. "Absolute" refers to the absolute value signs appearing in the definition.

As an example of the absolute distance, consider probability distributions on $\{0, 1\}$ defined by $p_0 = p, p_1 = 1-p$ and $q_0 = q, q_1 = 1-q$, where $p \leq q$. Then $D(p_x, q_x) = q-p+(1-p-(1-q)) = 2(q-p)$ is the absolute distance.

A second measure of distance between probability distributions, the *fidelity* of the probability distributions p_x and q_x is defined by

$$F(p_x, q_x) \equiv \sum_x \sqrt{p_x q_x}. \tag{5.2}$$

The fidelity is a quite different way of measuring distance between probability distributions than is the absolute distance. To begin with, it is not a metric, although we will see later that there is a metric which can be constructed from the fidelity. One way of seeing that the fidelity is not a metric

is to note that when the distributions p_x and q_x are identical, $F(p_x, q_x) = 1$. More generally, we will prove later that the fidelity is always in the range zero to one, and is equal to one if and only if the probability distributions are identical.

As an example of the fidelity, consider as before probability distributions on $0, 1$ defined by $p_0 = p, p_1 = 1 - p$ and $q_0 = q, q_1 = 1 - q$, where $p \leq q$. Then $F(p_x, q_x) = \sqrt{pq} + \sqrt{(1-p)(1-q)}$.

The absolute distance and fidelity are mathematically useful means of defining the notion of a distance between two probability distributions. Do these measures have physically motivated operational meanings? In the case of the absolute distance, the answer to this question is yes. In particular, it can be shown that [120]

$$D(p_x, q_x) = 2 \max_S \left(\sum_{x \in S} p_x - \sum_{x \in S} q_x) \right), \tag{5.3}$$

where the maximization is over all subsets S of the index set. The quantity being maximized is the difference between the probability that the event S occurs, according to the distribution p_x, and the probability that the event S occurs, according to the distribution q_x. The event S is thus the optimal event to examine when trying to distinguish the distributions p_x and q_x. The absolute distance governs how well it is possible to make this distinction, using statistical tools such as the Chernoff Bound [139].

I am not aware of a similarly clear operational interpretation for the fidelity. However, in the next section we will see that the fidelity is a sufficiently useful quantity for mathematical purposes to justify its study, even without a clear physical interpretation. Moreover, I can not rule out the possibility that a clear physical interpretation of the fidelity will be discovered in the future. Finally, it turns out that there are close connections between the fidelity and the absolute distance, which allow one to use properties of one quantity to deduce properties of the other.

The third notion of distance with which we are concerned is a *dynamical measure* of distance. Suppose the random variables X and Y form a Markov process[1], $X \rightarrow Y$, with values over the same possible range of values, which we denote by x. Then the probability that Y is not equal to X, $p(X \neq Y)$, is an obvious, but still important measure of the degree to which information has been preserved by the channel.

This measure of distance can be recast in the form of the absolute distance introduced earlier. Imagine that the random variable X is given to you, and you first make a copy of X, creating a new random variable $\tilde{X} = X$. The random variable X now undergoes some Markov dynamics, leaving as the output of the process a random variable, Y. How close is the initial perfectly correlated pair, $(\tilde{X}, X)$, to the final pair, $(\tilde{X}, Y)$? Using the absolute distance as our measure of distance, we see that the answer to this question is the absolute distance between the distributions $p_{x,x'} \equiv p(\tilde{X} = x, X = x') = \delta_{xx'} p(X = x)$ and $q_{x,x'} \equiv p(\tilde{X} = x, Y = x')$,

$$\begin{aligned}
D((\tilde{X}, X), (X, Y)) &= D(p_{x,x'}, q_{x,x'}) && (5.4)\\
&= \sum_{xx'} |\delta_{xx'} p(X = x) - p(\tilde{X} = x, Y = x')| && (5.5)\\
&= \sum_{x \neq x'} p(\tilde{X} = x, Y = x') + \sum_x \Big(p(X = x) - p(\tilde{X} = x, Y = x) \Big) && (5.6)\\
&= p(\tilde{X} \neq Y) + 1 - p(\tilde{X} = Y) && (5.7)\\
&= p(X \neq Y) + p(\tilde{X} \neq Y) && (5.8)
\end{aligned}$$

[1]Strictly speaking, any pair of random variables form a Markov process, but this usage is to get you in the mood for the less trivial Markov processes in the next paragraph.

$$= \quad 2p(X \neq Y). \tag{5.9}$$

It is worthwhile to reflect on this example. The probability of an error occurring during the Markov process is equal to (half) the absolute distance between $(\tilde{X}, X)$ and $(\tilde{X}, Y)$, which we can regard as a measure of the extent to which correlation between X and the external world is destroyed by the dynamics undergone by X. A quantum analogue of this idea will be used later to define a notion of information preservation through a quantum channel, based on the idea that it is quantum entanglement, rather than correlation, which is the important thing to preserve during the channel's dynamics.

The previous example concerned a Markov process containing only two random variables. More often, we will be concerned with a multipart Markov process. For example, imagine that we have a four part Markov process, $W \rightarrow X \rightarrow Y \rightarrow Z$. Such a situation arises, for example, in communications problems: W is an information source which is encoded using an error correcting code to give a random variable X, before being sent over a noisy communications channel which has Y as output, before being decoded to yield Z. Once again, the total probability of error, $p(W \neq Z)$ is an important distance measure for the channel.

5.2 How close are two quantum states?

What does it mean to say that two quantum states, ρ and σ, are close together? In this section we review two measures of the closeness of quantum states, the absolute distance and the fidelity, both of which generalize the corresponding classical concepts introduced in the previous section. Furthermore, we introduce two additional measures of distance, both of which arise naturally from the fidelity. The section concludes by examining relationships between the absolute distance and the fidelity.

5.2.1 Absolute distance

We begin by defining the *absolute distance* between states ρ and σ,

$$D(\rho, \sigma) \equiv \|\rho - \sigma\| \equiv \mathrm{tr}|\rho - \sigma|. \tag{5.10}$$

where $|A| \equiv \sqrt{A^\dagger A}$. Notice that this measure of distance generalizes the classical absolute distance, in the sense that if ρ and σ are diagonal in the same basis, then the (quantum) absolute distance between ρ and σ is equal to the classical absolute distance between the eigenvalues of ρ and σ.

There is a useful alternate formula for the absolute distance,

$$D(\rho, \sigma) = 2 \max_P \mathrm{tr}(P(\rho - \sigma)), \tag{5.11}$$

where the maximization may be taken alternately over all projectors, P, or over all positive operators $P \leq I$; the formula is valid in either case. This formula, which we shortly prove, gives rise to an appealing interpretation of the absolute distance. Using the identification of events which may occur as measurement outcomes in a quantum system with POVM elements – positive operators $P \leq I$ – we see that the absolute distance is equal to twice the difference in probabilities that an event P may occur, depending on whether the state is ρ or σ, maximized over all possible events P.

We prove this formula for the case where the maximization is over projectors; the case of positive operators $P \leq I$ follows the same reasoning. We begin by using the spectral decomposition of $\rho - \sigma$ to write $\rho - \sigma = Q - S$, where Q and S are positive operators with disjoint support.

Note that $|\rho - \sigma| = Q + S$, so $D(\rho, \sigma) = \mathrm{tr}(Q) + \mathrm{tr}(S)$. But $\mathrm{tr}(Q - S) = \mathrm{tr}(\rho - \sigma) = 0$, so $\mathrm{tr}(Q) = \mathrm{tr}(S)$, and therefore $D(\rho, \sigma) = 2\mathrm{tr}(Q)$. Let P be the projector onto the support of Q. Then $2\mathrm{tr}(P(\rho - \sigma)) = 2\mathrm{tr}(P(Q - S)) = 2\mathrm{tr}(Q) = D(\rho, \sigma)$. Conversely, let P be any projector. Then $2\mathrm{tr}(P(\rho - \sigma)) = 2\mathrm{tr}(P(Q - S)) \leq 2\mathrm{tr}(PQ) \leq 2\mathrm{tr}(Q) = D(\rho, \sigma)$. This completes the proof.

Perhaps the most important property of the absolute distance is that it is a metric on the space of density operators. It is clear that $D(\rho, \sigma) = 0$ if and only if $\rho = \sigma$, and that $D(\cdot, \cdot)$ is a symmetric function of its inputs. The triangle inequality,

$$D(\rho, \tau) \leq D(\rho, \sigma) + D(\sigma, \tau), \tag{5.12}$$

follows from the observation that there exist a projector P such that

$$\begin{aligned} D(\rho, \tau) &= 2\mathrm{tr}(P(\rho - \tau)) && (5.13)\\ &= 2\mathrm{tr}(P(\rho - \sigma)) + 2\mathrm{tr}(P(\sigma - \tau)) && (5.14)\\ &\leq D(\rho, \sigma) + D(\sigma, \tau). && (5.15) \end{aligned}$$

This completes the proof that the absolute distance is a metric.

The same method of proof can be used to show that the absolute distance is doubly convex in its inputs,

$$D(\sum_i p_i \rho_i, \sum_i p_i \sigma_i) \leq \sum_i p_i D(\rho_i, \sigma_i). \tag{5.16}$$

To see this, note that there exist a projector P such that

$$\begin{aligned} D(\sum_i p_i \rho_i, \sum_i p_i \sigma_i) &= 2\sum_i p_i \mathrm{tr}(P(\rho_i - \sigma_i)) && (5.17)\\ &\leq \sum_i p_i D(\rho_i, \sigma_i). && (5.18) \end{aligned}$$

Indeed, it is possible to prove a generalization of double convexity, using the same line of reasoning. Let p_i and q_i be probability distributions over the same index set. Then there exists a projector P such that

$$\begin{aligned} D(\sum_i p_i \rho_i, \sum_i q_i \sigma_i) &= 2\sum_i p_i \mathrm{tr}(P\rho_i) - 2\sum_i q_i \mathrm{tr}(P\sigma_i) && (5.19)\\ &= 2\sum_i p_i \mathrm{tr}(P(\rho_i - \sigma_i)) + 2\sum_i (p_i - q_i)\mathrm{tr}(P\sigma_i) && (5.20)\\ &\leq \sum_i p_i D(\rho_i, \sigma_i) + 2D(p_i, q_i), && (5.21) \end{aligned}$$

where $D(p_i, q_i)$ is the absolute distance between the probability distributions p_i and q_i.

Suppose $\mathcal{E}$ is a complete quantum operation. Let ρ and σ be density operators. Then Ruskai [149] has shown that

$$D(\mathcal{E}(\rho), \mathcal{E}(\sigma)) \leq D(\rho, \sigma). \tag{5.22}$$

That is, physical quantum operations are contractive maps on the space of density operators. To prove this, use the spectral decomposition to write $\rho - \sigma = Q - S$, where Q and S are positive matrices

with disjoint support, and let P be a projector such that $D(\mathcal{E}(\rho), \mathcal{E}(\sigma)) = 2\mathrm{tr}(P(\mathcal{E}(\rho) - \mathcal{E}(\sigma)))$. Note that

$$\begin{aligned} D(\rho, \sigma) &= \mathrm{tr}|Q - S| && (5.23)\\ &= \mathrm{tr}(Q) + \mathrm{tr}(S) && (5.24)\\ &= \mathrm{tr}(\mathcal{E}(Q)) + \mathrm{tr}(\mathcal{E}(S)) && (5.25)\\ &= 2\mathrm{tr}(\mathcal{E}(Q)) && (5.26)\\ &\geq 2\mathrm{tr}(P\mathcal{E}(Q)) && (5.27)\\ &\geq 2\mathrm{tr}(P(\mathcal{E}(Q) - \mathcal{E}(S))) && (5.28)\\ &= 2\mathrm{tr}(P(\mathcal{E}(\rho) - \mathcal{E}(\sigma))) && (5.29)\\ &= D(\mathcal{E}(\rho), \mathcal{E}(\sigma)), && (5.30) \end{aligned}$$

which completes the proof.

Contractivity together with double convexity can be used to prove results about the existence of stationary states for a quantum operation. Suppose $\mathcal{E}$ is a quantum operation for which there exists a fixed density operator ρ_0 and a quantum operation $\mathcal{E}'$ such that

$$\mathcal{E}(\rho) = p\rho_0 + (1-p)\mathcal{E}'(\rho), \tag{5.31}$$

for some p, $0 < p \leq 1$. Physically, this means that with a certain probability p, the input state is thrown out and replaced with the fixed state ρ_0. With probability $1-p$, the operation $\mathcal{E}'$ is applied. An important example of a channel of this type is the much-studied *depolarizing channel* for a qubit [22], which with probability p randomizes the state, that is, replaces it with the fixed operator $I/2$, and with probability $1-p$ leaves the state untouched. By the double convexity of the absolute distance, it follows that

$$\begin{aligned} D(\mathcal{E}(\rho), \mathcal{E}(\sigma)) &\leq pD(\rho_0, \rho_0) + (1-p)D(\mathcal{E}'(\rho), \mathcal{E}'(\sigma)) && (5.32)\\ &\leq (1-p)D(\rho, \sigma), && (5.33) \end{aligned}$$

where on the second line we have applied the contractivity of the absolute distance with respect to physical quantum operations. Thus, the class of quantum operations which have this form are *strictly contractive*, and it is not difficult to see that they have a unique fixed point; see the Lemma in Appendix One of [167] for a proof.

We noticed earlier that the absolute distance has an interpretation as half the maximal difference in probabilities that may arise from a single measurement result on the two density operators. We now explore a slightly different way of viewing the operational meaning of the absolute distance. Suppose M_m is a set of POVM elements describing a measurement on the quantum system. Let $p_m \equiv \mathrm{tr}(\rho M_m)$ and $q_m \equiv \mathrm{tr}(\sigma M_m)$ be the probabilities associated with the POVM measurement. Then $D(p_m, q_m) \leq D(\rho, \sigma)$. To see this, note that

$$D(p_m, q_m) = \sum_m |\mathrm{tr}(M_m(\rho - \sigma))|. \tag{5.34}$$

Using the spectral theorem we may decompose $\rho - \sigma = Q - S$, where Q and S are positive operators with disjoint support. Thus $|\rho - \sigma| = Q + S$, and

$$\begin{aligned} |\mathrm{tr}(M_m(\rho - \sigma))| &= |\mathrm{tr}(M_m(Q - S))| && (5.35)\\ &\leq \mathrm{tr}(M_m(Q + S)) && (5.36)\\ &\leq \mathrm{tr}(M_m|\rho - \sigma|). && (5.37) \end{aligned}$$

Thus

$$\begin{aligned} D(p_m, q_m) &\leq \sum_m \mathrm{tr}(M_m|\rho - \sigma|) && (5.38)\\ &= \mathrm{tr}(|\rho - \sigma|) && (5.39)\\ &= D(\rho, \sigma), && (5.40) \end{aligned}$$

where we have applied the completeness relation for POVM elements, $\sum_m M_m = I$.

Thus, if two density operators are close in absolute distance, then any measurement performed on those quantum states will give rise to probability distributions which are close together in the classical sense of absolute distance. Conversely, by choosing a measurement whose POVM elements include projectors onto the support of Q and S, we see that there exist measurements which give rise to probability distributions such that $D(p_m, q_m) = D(\rho, \sigma)$.

Thus, we have a second interpretation of the absolute distance between two quantum states, as an achievable upper bound on the absolute distance between probability distributions arising from measurements performed on those quantum states.

We conclude our survey of elementary properties of the absolute distance with an elegant result linking the absolute distance to entropy. This result is known as *Fannes' inequality* [138]. It states that for density operators ρ and σ such that $D(\rho, \sigma) \leq 1/e$,

$$|S(\rho) - S(\sigma)| \leq D(\rho, \sigma) \log d + \eta(D(\rho, \sigma)), \tag{5.41}$$

where d is the dimensionality of the underlying Hilbert space, and $\eta(x) \equiv -x \log x$.

To prove Fannes' inequality we need a simple result relating the absolute distance between two operators to their eigenvalues. Let $r_1 \geq r_2 \geq \ldots \geq r_d$ be the eigenvalues of ρ, in descending order, with corresponding orthonormal eigenvectors $|e_i\rangle$, and $s_1 \geq s_2 \geq \ldots \geq s_d$ be the eigenvalues of σ, again in descending order, with corresponding eigenvectors $|f_i\rangle$. Then decompose $\rho - \sigma = Q - R$, where Q and R are positive operators with disjoint support. Defining $T \equiv R + \rho = Q + \sigma$, we have

$$D(\rho, \sigma) = \mathrm{tr}(R + Q) = \mathrm{tr}(2T) - \mathrm{tr}(R) - \mathrm{tr}(Q). \tag{5.42}$$

Let $t_1 \geq t_2 \geq \ldots \geq t_d$ be the eigenvalues of T. Note that $t_i \geq \max(r_i, s_i)$, so $2t_i \geq r_i + s_i + |r_i - s_i|$. From equation (5.42) it follows that

$$D(\rho, \sigma) \geq \sum_i |r_i - s_i|, \tag{5.43}$$

which is the relation we shall need to prove Fannes' inequality.

To prove Fannes' inequality we use the inequality

$$|\eta(r) - \eta(s)| \leq \eta(|r - s|), \tag{5.44}$$

which may be easily verified by calculus whenever $|r - s| \leq 1/2$. A little thought shows that $|r_i - s_i| \leq 1/2$ for all i, so

$$\begin{aligned} |S(\rho) - S(\sigma)| &= \left| \sum_i (\eta(r_i) - \eta(s_i)) \right| && (5.45)\\ &\leq \sum_i \eta(|r_i - s_i|). && (5.46) \end{aligned}$$

Setting $\Delta \equiv \sum_i |r_i - s_i|$, we see that

$$\begin{aligned} |S(\rho) - S(\sigma)| &\leq \Delta\eta(|r_i - s_i|/\Delta) + \eta(\Delta) && (5.47) \\ &\leq \Delta \log d + \eta(\Delta). && (5.48) \end{aligned}$$

But $\Delta \leq D(\rho, \sigma)$ by (5.43), so by the monotonicity of η on the interval $[0, 1/e]$,

$$|S(\rho) - S(\sigma)| \leq D(\rho, \sigma) \log d + \eta(D(\rho, \sigma)), \tag{5.49}$$

whenever $D(\rho, \sigma) \leq 1/e$, which is Fannes' inequality. A minor modification to the previous reasoning shows that for general $D(\rho, \sigma)$, the slightly weaker form of Fannes' inequality,

$$|S(\rho) - S(\sigma)| \leq D(\rho, \sigma) \log d + \frac{1}{e}, \tag{5.50}$$

holds.

5.2.2 Fidelity

A second measure of distance between two quantum states is the *fidelity*. This subsection reviews the definition and basic properties of the fidelity. At the outset, it is well to mind that the fidelity is not a true measure of distance, as it is not a metric, but it does give rise to a metric, which will be reviewed in the next subsection.

The fidelity of states ρ and σ is defined to be[2] [178, 92, 68]

$$F(\rho, \sigma) \equiv \operatorname{tr} \sqrt{\rho^{1/2} \sigma \rho^{1/2}}. \tag{5.51}$$

Note that when ρ and σ are diagonal in the same basis, this reduces to the classical fidelity between the eigenvalues of the two states.

There is a useful alternative characterization of the fidelity due to Uhlmann [178]. Suppose we denote the quantum system where our states live by the letter Q. Introduce another quantum system, R, which is a copy of Q. Then, as discussed in Appendix A, for any mixed state ρ of Q, it is possible to find a pure state $|\psi\rangle$ of RQ such that $|\psi\rangle$ extends ρ in the natural way. We call such a $|\psi\rangle$ a *purification* of ρ. It can be shown that

$$F(\rho, \sigma) = \max_{|\psi\rangle, |\phi\rangle} |\langle \psi | \phi \rangle|, \tag{5.52}$$

where the maximization is performed over all purifications $|\psi\rangle$ of ρ, and $|\phi\rangle$ of σ. We will not prove this formula here, but instead refer the reader to [68] for an elegant proof. There are several variants of this formula which are easily seen to be equivalent. For instance, it is possible to fix any purification $|\psi\rangle$ of ρ, and simply maximize over purifications of σ. Moreover, purifying the states ρ and σ into the space RQ was not necessary; any space large enough to contain purifications of both ρ and σ will suffice.

Uhlmann's formula does not provide a calculational tool for evaluating the fidelity, as does equation (5.51). However, in many instances, properties of the fidelity are more easily proved using Uhlmann's formula than equation (5.51).

Uhlmann's formula makes it clear that the fidelity is symmetric in its inputs, $F(\rho, \sigma) = F(\sigma, \rho)$, and that the fidelity is bounded between 0 and 1, $0 \leq F(\rho, \sigma) \leq 1$. If $\rho = \sigma$ then it is clear

[2] The reader ought to be aware that in the literature both the quantity we call fidelity and its square have been referred to as the fidelity. Compare also the definition of the dynamic fidelity given below, in section 5.3.

that $F(\rho, \sigma) = 1$, from Uhlmann's formula. If $\rho \neq \sigma$ then $|\psi\rangle \neq |\phi\rangle$ for any purifications $|\psi\rangle$ and $|\phi\rangle$ of ρ and σ, respectively, so $F(\rho, \sigma) < 1$. From equation (5.51) we see that $F(\rho, \sigma) = 0$ if and only if ρ and σ have disjoint support.

Summarizing, the fidelity is symmetric in its inputs, $0 \leq F(\rho, \sigma) \leq 1$, with equality in the first if and only if ρ and σ have orthogonal support, and equality in the second if and only if $\rho = \sigma$.

There is a simple instance in which a useful explicit formula for the fidelity may be given. Suppose we wish to calculate the fidelity between a pure state $|\psi\rangle$ and an arbitrary state, ρ. From equation (5.51) we see that

$$\begin{aligned} F(|\psi\rangle, \rho) &= \mathrm{tr}\sqrt{\langle\psi|\rho|\psi\rangle \, |\psi\rangle\langle\psi|} && (5.53) \\ &= \sqrt{\langle\psi|\rho|\psi\rangle}. && (5.54) \end{aligned}$$

That is, the fidelity is equal to the square root of the overlap between $|\psi\rangle$ and ρ. This is an important result which we will make much use of.

As already noted, the fidelity is not a metric. However, in many other ways the fidelity closely resembles the absolute distance. The remainder of this section is used to prove two results about the fidelity which are analogous to properties already proved of the absolute distance. These results concern, respectively, a strong concavity result for the fidelity; and a proof that the fidelity can not increase under quantum operations.

First, we examine the concavity properties of the fidelity. We will use the Uhlmann formula for fidelity to prove a strong concavity property for the fidelity. Let p_i and q_i be probability distributions over the same index set, and ρ_i and σ_i density operators also indexed by the same index set. Then

$$F(\sum_i p_i \rho_i, \sum_i q_i \sigma_i) \geq \sum_i \sqrt{p_i q_i} F(\rho_i, \sigma_i). \quad (5.55)$$

To see this, let $|\psi_i\rangle$ and $|\phi_i\rangle$ be purifications of ρ_i and σ_i chosen such that $F(\rho_i, \sigma_i) = \langle\psi_i|\phi_i\rangle$. Introduce a system I which has orthonormal basis states $|i\rangle$ corresponding to the index set i for the probability distributions. Define

$$\begin{aligned} |\psi\rangle &\equiv \sum_i \sqrt{p_i} |\psi_i\rangle |i\rangle && (5.56) \\ |\phi\rangle &\equiv \sum_i \sqrt{q_i} |\phi_i\rangle |i\rangle. && (5.57) \end{aligned}$$

Note that $|\psi\rangle$ is a purification of $\sum_i p_i \rho_i$ and $|\phi\rangle$ is a purification of $\sum_i q_i \sigma_i$, so by Uhlmann's formula,

$$\begin{aligned} F(\sum_i p_i \rho_i, \sum_i q_i \sigma_i) &\geq |\langle\psi|\phi\rangle| && (5.58) \\ &= \sum_i \sqrt{p_i q_i} \langle\psi_i|\phi_i\rangle && (5.59) \\ &= \sum_i \sqrt{p_i q_i} F(\rho_i, \sigma_i), && (5.60) \end{aligned}$$

which establishes the result we set out to prove. We refer to this result as the *strong concavity* of the fidelity.

The strong concavity of the fidelity has a number of useful consequences. One is the *joint concavity* of the fidelity. In particular, note that if $p_i = q_i$, then strong concavity reduces to

$$F(\sum_i p_i \rho_i, \sum_i p_i \sigma_i) \geq \sum_i p_i F(\rho_i, \sigma_i). \tag{5.61}$$

The joint concavity, in turn, implies that the fidelity is concave in each entry. For example, for each i set $\rho_i = \rho$ for some fixed ρ. Then the joint concavity of the fidelity reduces to

$$F(\rho, \sum_i p_i \sigma_i) \geq \sum_i p_i F(\rho, \sigma_i), \tag{5.62}$$

that is, the fidelity is concave in the second entry. By symmetry, the fidelity is also concave in the first entry.

The second property of the fidelity which we prove is that it is non-decreasing under complete quantum operations [8],

$$F(\mathcal{E}(\rho), \mathcal{E}(\sigma)) \geq F(\rho, \sigma). \tag{5.63}$$

To prove this, let $|\psi\rangle$ and $|\phi\rangle$ be purifications of ρ and σ into a joint system RQ such that $F(\rho, \sigma) = |\langle\psi|\phi\rangle|$. Introduce a model environment E for the quantum operation, $\mathcal{E}$, which starts in a pure state $|0\rangle$, and interacts with the quantum system Q via a unitary interaction U. Note that $U|\psi\rangle|0\rangle$ is a purification of $\mathcal{E}(\rho)$, and $U|\phi\rangle|0\rangle$ is a purification of $\mathcal{E}(\sigma)$. By Uhlmann's formula it follows that

$$\begin{aligned} F(\mathcal{E}(\rho), \mathcal{E}(\sigma)) &\geq |\langle\psi|\langle 0|U^\dagger U|\phi\rangle|0\rangle| && (5.64)\\ &= |\langle\psi|\phi\rangle| && (5.65)\\ &= F(\rho, \sigma), && (5.66) \end{aligned}$$

establishing the property that we set out to prove.

This completes our discussion of elementary properties of the fidelity. Note that the fidelity and the absolute distance have many similar properties, although I am not aware of any simple physical interpretation of the fidelity. Why do we bother developing both quantities? We do so because it often helps to have more than one way of doing things; one obtains new insights from multiple ways of viewing the same phenomena. It is also potentially the case that in the future a powerful property of one of these quantities will be found that has no natural analogue which applies to the other quantity. Indeed, I have seen fit to discuss both the fidelity and the absolute distance in this Chapter because most of the research later in the Dissertation has been carried out using the fidelity as a tool, while now it seems to me that the absolute distance has a more compelling physical interpretation, and is equally powerful mathematically.

5.2.3 Distance measures derived from fidelity

The fidelity may be used to develop many other useful measures of distance between density operators. This subsection develops two natural measures of distance derived from the fidelity, the *error* and the *angle*, and develops some elementary properties of these measures, most importantly, the fact that the angle is a metric on the space of density operators.

Given that the fidelity is bounded between 0 and 1, and is equal to one if and only if the states being compared are equal, the most obvious candidate for a metric is the function defined by

$$\tilde{E}(\rho, \sigma) \equiv 1 - F(\rho, \sigma). \tag{5.67}$$

It turns out, however, that it is slightly more convenient to define the function

$$E(\rho, \sigma) \equiv 1 - F(\rho, \sigma)^2, \tag{5.68}$$

which we shall refer to as the *error* for ρ and σ. Both functions are metrics on density operators, and have many other nice properties, however it turns out that the error function has properties that will be of especial use in the study of dynamic measures of distance, properties which $\tilde{E}$ does not have. The error has numerous useful properties which it inherits from the fidelity:

1. $E(\rho, \sigma) = 0$ if and only if $\rho = \sigma$.

2. Symmetry. $E(\rho, \sigma) = E(\sigma, \rho)$.

3. Let $\mathcal{E}$ be a complete quantum operation. Then

$$E(\mathcal{E}(\rho), \mathcal{E}(\sigma)) \leq E(\rho, \sigma). \tag{5.69}$$

The error will assume a significant role in later discussions of dynamic measures of distance, and throughout the remainder of this Dissertation. We now switch to a second measure of distance derived from the fidelity, the *angle*. This measure will play a much lesser role in the remainder of this Dissertation. It is included here as a teaser to indicate just one of the wide variety of natural directions which research into distance measures for quantum states may take.

Recall Uhlmann's formula, that the fidelity between two states is equal to the maximum inner product between purifications of those states. Recall that in Cartesian geometry the inner product between two unit vectors has an interpretation as the cosine of the angle between these states. This suggests that we define the *generalized angle* between states ρ and σ by

$$A(\rho, \sigma) \equiv \arccos F(\rho, \sigma). \tag{5.70}$$

The generalized angle, which we will usually refer to just as the angle, is a real number in the range 0 to $\pi/2$. The angle is also a true distance measure on density operators. The following is a summary of the elementary properties of the angle, each of which is immediate from properties of the fidelity, together with the observations from calculus that arccos is a decreasing concave function on the interval $[0, 1]$. Where this is not the case a brief proof is given.

1. $A(\rho, \sigma) = 0$ if and only if $\rho = \sigma$.

2. Symmetry. $A(\rho, \sigma) = A(\sigma, \rho)$.

3. $A(\rho, \sigma)$ satisfies the triangle inequality, and therefore is a metric. This is immediate from Uhlmann's formula, and the definition of the angle.

4. Let $\mathcal{E}$ be a complete quantum operation. Then

$$A(\mathcal{E}(\rho), \mathcal{E}(\sigma)) \leq A(\rho, \sigma). \tag{5.71}$$

5.2.4 Relationships between distance measures

There are several useful relationships between the absolute distance and the fidelity.

Consider the absolute distance between two pure states, $|a\rangle$ and $|b\rangle$. Introduce orthonormal states $|0\rangle$ and $|1\rangle$ such that $|a\rangle = |0\rangle$ and $|b\rangle = \cos\theta|0\rangle + \sin\theta|1\rangle$. Notice that $F(|a\rangle, |b\rangle) = |\cos\theta|$. Furthermore,

$$\begin{aligned} D(|a\rangle, |b\rangle) &= \mathrm{tr}\left|\begin{bmatrix} 1-\cos^2\theta & -\cos\theta\sin\theta \\ -\cos\theta\sin\theta & \sin^2\theta \end{bmatrix}\right| && (5.72) \\ &= 2|\sin\theta| && (5.73) \\ &= 2\sqrt{1-F(|a\rangle, |b\rangle)^2} = 2\sqrt{E(|a\rangle, |b\rangle)}. && (5.74) \end{aligned}$$

Let ρ and σ be any two quantum states, and let $|\psi\rangle$ and $|\phi\rangle$ be purifications chosen such that $F(\rho, \sigma) = |\langle\psi|\phi\rangle| = F(|\psi\rangle, |\phi\rangle)$. Recalling that absolute distance is non-increasing under the partial trace, we see that

$$\begin{aligned} D(\rho, \sigma) &\leq D(|\psi\rangle, |\phi\rangle) && (5.75) \\ &= 2\sqrt{1-F(\rho,\sigma)^2} = 2\sqrt{E(\rho,\sigma)}. && (5.76) \end{aligned}$$

Thus, if the error between two states is small, it will follow that the states are also close in absolute distance. The converse is also true, at least when one of the two states is a pure state, which will be sufficient for the applications we shall consider. Let $|\psi\rangle$ be a pure state, and σ an arbitrary state. Then

$$\begin{aligned} D(|\psi\rangle, \sigma) &= 2\max_P \mathrm{tr}(P(|\psi\rangle\langle\psi| - \sigma)) && (5.77) \\ &\geq 2\mathrm{tr}(|\psi\rangle\langle\psi|(|\psi\rangle\langle\psi| - \sigma)) && (5.78) \\ &= 2(1 - F(|\psi\rangle, \sigma)^2). && (5.79) \end{aligned}$$

Restating these bounds in term of the error, we see that

$$2E(|\psi\rangle, \sigma) \leq D(|\psi\rangle, \sigma)) \leq 2\sqrt{E(|\psi\rangle, \sigma)}. \quad (5.80)$$

The implication of this relation is that when one of the inputs is a pure state, and the other state is arbitrary, the absolute distance and the error are equally good measures of closeness for quantum states, at least in terms of their limiting behaviours. Part II of this Dissertation is largely concerned with such limiting behaviours; in such instances this relation implies that it does not matter whether the error or the absolute distance is used as a measure of distance, since any result about one will imply a qualitatively similar result about the other.

5.3 Dynamic measures of information preservation

This section uses the static measures of distance discussed in previous sections to develop several measures of how well a quantum operation preserves information. A major concern of this Dissertation is the transmission of entangled states through quantum channels, so we will focus on measures related to this problem.

We will primarily be interested in the following model scenario. A quantum system, Q, is prepared in a state ρ. The state of Q is entangled in some way with the external world. We represent this entanglement by introducing a fictitious system R, such that the joint state of RQ is a pure state. It turns out that all results that we prove do not depend in any way on how this purification is performed, so we may as well suppose that this is an arbitrary entanglement with the outside

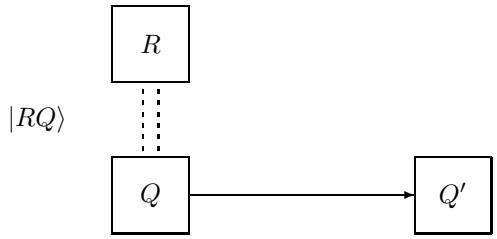

Figure 5.1: The RQ picture of a quantum channel. The initial state of RQ is a pure state.

world. The system Q is then subjected to a dynamics described by a quantum operation, $\mathcal{E}$. The basic situation is illustrated in figure 5.1.

How well is the entanglement between R and Q preserved by the quantum operation $\mathcal{E}$? We investigate two ways of quantifying this. The first measure we refer to as the *dynamic distance*. It is defined by the expression

$$D(\rho, \mathcal{E}) \equiv D(RQ, R'Q'), \tag{5.81}$$

where the use of a prime indicates the state of a system after the quantum operation has been applied, and the absence of a prime indicates the state of a system before the quantum operation has been applied. Note that this expression depends only upon ρ and $\mathcal{E}$, and not upon the details of the purification RQ. To see this, we use the fact, proved in Appendix A, that any two purifications R_1Q_1 and R_2Q_2 of ρ are related by a unitary operation, U, *that acts upon* R *alone*, $R_2Q_2 = U(R_1Q_1)U^\dagger$. Thus

$$\begin{aligned} D(R_2Q_2, R_2'Q_2') &= \mathrm{tr}|R_2Q_2 - R_2'Q_2'| && (5.82)\\ &= \mathrm{tr}(U|R_1Q_1 - R_1'Q_1'|U^\dagger) && (5.83)\\ &= \mathrm{tr}|R_1Q_1 - R_1'Q_1'| && (5.84)\\ &= D(R_1Q_1, R_1'Q_2'), && (5.85) \end{aligned}$$

which establishes the result. Notice that the dynamic distance provides a measure of how well the entanglement between Q and R is preserved by the process, with values close to zero indicating that the entanglement has been very well preserved, and larger values indicating that it has not been so well preserved.

The second measure[3] quantifying how well the entanglement is preserved is known as the *entanglement fidelity*, although we will usually refer to it just as the *fidelity*, or the *dynamic fidelity*. As for the dynamic distance, the dynamic fidelity is defined for a *process*, specified by a quantum operation $\mathcal{E}$ acting on some initial state, ρ. We denote it by $F(\rho, \mathcal{E})$.

Formally, the dynamic fidelity is defined by

$$\begin{aligned} F(\rho, \mathcal{E}) &\equiv F(RQ, R'Q')^2 && (5.86)\\ &= \frac{\langle RQ|(\mathcal{I}_R \otimes \mathcal{E})(|RQ\rangle\langle RQ|)|RQ\rangle}{\mathrm{tr}(\mathcal{I}_R \otimes \mathcal{E})(|RQ\rangle\langle RQ|)}, && (5.87) \end{aligned}$$

[3]Historically, this measure was introduced earlier than the dynamic distance. It was introduced by Schumacher [157].

where the quantity appearing on the right hand side is the static fidelity between the initial and final states of RQ, *squared*. Thus, the dynamic fidelity provides a measure of how well the entanglement between R and Q is preserved by the process $\mathcal{E}$, with values close to 1 indicating that the entanglement has been well preserved, and values close to 0 indicating that most of the entanglement has been destroyed. The choice of whether to use the static fidelity squared or the static fidelity is essentially arbitrary; the present definition seems to result in slightly more attractive mathematical properties. Recall that the fidelity is not known to have a clear physical interpretation, so in this instance it is legitimate for us to make a decision based upon mathematical elegance, rather than physical necessity. To go with the dynamic fidelity, we define the *dynamic error* in a manner analogous to the earlier definition of the static error,

$$E(\rho, \mathcal{E}) \equiv 1 - F(\rho, \mathcal{E}). \tag{5.88}$$

Note that there is no square appearing on the right hand side, since the dynamic fidelity is already a fidelity squared. Thus,

$$E(\rho, \mathcal{E}) = E(RQ, R'Q'), \tag{5.89}$$

where the $E(\cdot, \cdot)$ appearing on the right hand side is the static error introduced in the previous section.

Note further that the dynamic fidelity and dynamic error do not depend upon the particular purification RQ of Q that is chosen, but only upon the initial state Q, and the quantum operation $\mathcal{E}$ [157]. The proof of this is as for the proof of the dynamic distance, which was adapted from [157].

Suppose E_i is a set of operation elements for a quantum operation $\mathcal{E}$. Then

$$F(\rho, \mathcal{E}) = \langle RQ|R'Q'|RQ\rangle = \frac{\sum_i |\langle RQ|E_i|RQ\rangle|^2}{\mathrm{tr}(\mathcal{E}(\rho))}. \tag{5.90}$$

Note that

$$\begin{aligned} \langle RQ|E_i|RQ\rangle &= \sum_{jk} \sqrt{p_j p_k}\langle j|k\rangle\langle j|E_i|k\rangle & (5.91)\\ &= \sum_j p_j \langle j|E_i|j\rangle & (5.92)\\ &= \mathrm{tr}(E_i\rho). & (5.93) \end{aligned}$$

Combining this expression with equation (5.90) we obtain the useful computational formula [157, 134]

$$F(\rho, \mathcal{E}) = \frac{\sum_i |\mathrm{tr}(\rho E_i)|^2}{\mathrm{tr}(\mathcal{E}(\rho))}. \tag{5.94}$$

This expression simplifies for trace-preserving quantum operations since the denominator is one. The dynamic fidelity has the following properties, with corresponding properties of the dynamic error obvious corollaries:

1. $0 \leq F(\rho, \mathcal{E}) \leq 1$ [157]. Follows from properties of the static fidelity.

2. The dynamic fidelity is convex in the density operator input, and linear in the quantum operation input, for complete quantum operations [157, 10]. The linearity is immediate from the definition of the dynamic fidelity. The convexity may be proved in many ways; one simple

technique is to use equation (5.94) to show that the function $f(x) \equiv F(x\rho_1 + (1-x)\rho_2, \mathcal{E})$ has a non-negative second derivative. Elementary calculus shows that

$$f''(x) = \sum_i |\mathrm{tr}((\rho_1 - \rho_2)E_i)|^2 \geq 0, \tag{5.95}$$

as required.

3. Unfortunately, the dynamic fidelity is not jointly convex in the density operator and the quantum operation. To see this, consider the following example on a space spanned by orthonormal states $|0\rangle, |1\rangle, |2\rangle, |3\rangle$. Let P_0, P_1, P_2, P_3 denote projections onto the corresponding one dimensional subspaces, and P_{12} and P_{34} projections onto the corresponding two dimensional subspaces. Define complete quantum operations $\mathcal{E}_1$ and $\mathcal{E}_2$ by $\mathcal{E}_1(\rho) \equiv P_1\rho P_1 + P_2\rho P_2 + P_{34}\rho P_{34}$ and $\mathcal{E}_2(\rho) \equiv P_{12}\rho P_{12} + P_3\rho P_3 + P_4\rho P_4$. Then

$$F\left(\frac{I}{4}, \frac{\mathcal{E}_1}{2} + \frac{\mathcal{E}_2}{2}\right) = \frac{3}{8} \not\leq \frac{1}{4} = \frac{1}{2}F\left(\frac{P_{12}}{2}, \mathcal{E}_1\right) + \frac{1}{2}F\left(\frac{P_{34}}{2}, \mathcal{E}_2\right), \tag{5.96}$$

so this is the desired counterexample.

4. The dynamic fidelity is a lower bound on the static fidelity squared between the input and output to the process [157],

$$F(\rho, \mathcal{E}) \leq \left[F(\rho, \mathcal{E}(\rho)/\mathrm{tr}(\mathcal{E}(\rho)))\right]^2. \tag{5.97}$$

The proof is an elementary application of the non-decreasing property of the static fidelity under partial trace, $F(\rho, \mathcal{E}) = F(RQ, R'Q')^2 \leq F(Q, Q')^2$. The intuitive meaning of the result is obviously a desirable property: it is harder to preserve a state plus entanglement with the outside world than it is to merely preserve the state alone.

5. For pure state inputs, the dynamic fidelity is equal to the static fidelity squared between input and output,

$$F(|\psi\rangle, \mathcal{E}) = F(|\psi\rangle, \mathcal{E}(|\psi\rangle\langle\psi|)/\mathrm{tr}(\mathcal{E}(|\psi\rangle\langle\psi|)))^2. \tag{5.98}$$

This is immediate from the observation that the state $|\psi\rangle$ is a purification of itself, and the definition.

6. Using convexity and the result that the dynamic fidelity is a lower bound on the static fidelity squared, we see that if $\{p_i, \rho_i\}$ is an ensemble of states generating ρ then

$$F(\rho, \mathcal{E}) \leq \sum_i p_i F(\rho_i, \mathcal{E}(\rho_i))^2, \tag{5.99}$$

for a complete quantum operation $\mathcal{E}$.

7. $F(\rho, \mathcal{E}) = 1$ if and only if for all pure states $|\psi\rangle$ lying in the support of ρ,

$$\mathcal{E}(|\psi\rangle\langle\psi|) = |\psi\rangle\langle\psi|. \tag{5.100}$$

Suppose $F(\rho, \mathcal{E}) = 1$, and $|\psi\rangle$ is a pure state in the support of ρ. Define $p \equiv \langle\psi|\rho|\psi\rangle > 0$ and σ to be a density operator such that $(1-p)\sigma = \rho - p|\psi\rangle\langle\psi|$. Then by convexity,

$$1 = F(\rho, \mathcal{E}) \leq pF(|\psi\rangle, \mathcal{E}) + (1-p), \tag{5.101}$$

and thus $F(|\psi\rangle, \mathcal{E}) = 1$, establishing the result one way. The other way is a straightforward application of the definition of the dynamic fidelity. This result was proved in [158], via a different technique, based upon the next property in this list.

8. The following result, due to Knill and Laflamme [98], is essentially a strengthening of the previous result. Suppose that $\langle\psi|\mathcal{E}(|\psi\rangle\langle\psi|)|\psi\rangle \geq 1 - \eta$ for all $|\psi\rangle$ in the support of ρ, for some η. Then $F(\rho, \mathcal{E}) \geq 1 - (3\eta/2)$.

Earlier we derived equation (5.80) which related the absolute distance and the fidelity in the static case, when one of the inputs was a pure state. This set of inequalities immediately implies a corresponding set of inequalities for the dynamic distance and dynamic error,

$$2E(\rho, \mathcal{E}) \leq D(\rho, \mathcal{E}) \leq 2\sqrt{E(\rho, \mathcal{E})}. \tag{5.102}$$

This is a very useful result. First, it tells us that the dynamic distance and dynamic error are essentially equivalent as measures of how well entanglement is preserved when it undergoes a process, that is, the dynamic distance is small if and only if the dynamic error is small. Second, the result gives us alternate means for investigating quantum channels: if the dynamic error, for instance, is proving to be difficult to use, one can switch to the dynamic distance in an effort to simplify the analysis. If successful, the results obtained using the dynamic distance can then be translated back in terms of the dynamic error. Two measures are better than one.

5.3.1 Continuity relations

What continuity properties are possessed by the dynamic measures of information preservation? Naturally, we expect that if the input to a quantum process is perturbed slightly, then the distance measures associated with that process should only change by a small amount. In this section we give a bound on the extent to which this is true.

We will call such relations *continuity relations*, although this is perhaps a slightly dubious coinage. After all, the distance measures being investigated are defined in terms of the self-same metrics with respect to which we are investigating their continuity properties.

By the triangle inequality and the non-increasing property of the absolute distance under quantum operations,

$$\begin{aligned} D(\rho_1, \mathcal{E}) &= D(R_1 Q_1, R_1' Q_1') && (5.103)\\ &\leq D(R_1 Q_1, R_2 Q_2) + D(R_2 Q_2, R_2' Q_2') + D(R_2' Q_2', R_1' Q_1') && (5.104)\\ &\leq 2D(R_1 Q_1, R_2 Q_2) + D(\rho_2, \mathcal{E}). && (5.105) \end{aligned}$$

Minimizing over purifications of ρ_1 and ρ_2 we obtain

$$D(\rho_1, \mathcal{E}) \leq D(\rho_2, \mathcal{E}) + 4\sqrt{E(\rho_1, \rho_2)}. \tag{5.106}$$

Thus, if ρ_1 is close to ρ_2, as measured by the fidelity, then $D(\rho_1, \mathcal{E})$ and $D(\rho_2, \mathcal{E})$ must also be close together.

5.3.2 Chaining quantum errors

Suppose we have a composite quantum process generated by quantum operations $\mathcal{E}_1$ and $\mathcal{E}_2$,

$$\rho \xrightarrow{\mathcal{E}_1} \rho' \xrightarrow{\mathcal{E}_2} \rho''. \tag{5.107}$$

Is there a way of relating the absolute distance associated with the complete two-part process with the absolute distance associated with the component processes? We will show that

$$D(\rho, \mathcal{E}_2 \circ \mathcal{E}_1) \leq D(\rho, \mathcal{E}_1) + D(\rho', \mathcal{E}_2). \tag{5.108}$$

Thus, in order for there to be little error in the composite process $\mathcal{E}_2 \circ \mathcal{E}_2$, it suffices that there be little error caused by the process $\mathcal{E}_1$ or $\mathcal{E}_2$.

To see this, introduce a system R_2' which purifies the system $R'Q'$. Then notice that

$$D(R'Q', R''Q'') \leq D(R'R_2'Q', R''R_2''Q'') = D(\rho', \mathcal{E}_2). \tag{5.109}$$

Thus

$$\begin{aligned} D(\rho, \mathcal{E}_2 \circ \mathcal{E}_1) &= D(RQ, R''Q'') && (5.110)\\ &\leq D(RQ, R'Q') + D(R'Q', R''Q'') && (5.111)\\ &\leq D(\rho, \mathcal{E}_1) + D(\rho', \mathcal{E}_2), && (5.112) \end{aligned}$$

as we set out to demonstrate.

This result about the chaining behaviour of errors will not be used much later in the Dissertation. Nevertheless, it is very important conceptually. Essentially, it tells us in a quantitative way that if we want to ensure that a complicated quantum process is is carried out well, then it is sufficient to ensure that each step of that process is carried out reliably.

5.4 Alternative view of the dynamic measures

In the previous section we presented several dynamic measures of quantum information preservation based upon the RQ picture of a quantum process. In this picture, the state of a quantum system, Q, is first purified into a fictitious quantum system, R. R is used to represent the possibility that Q is entangled with another system.

In this section we will give a more obviously physical account of the dynamic measures of information. To avoid repetition, this account will be phrased entirely in terms of the dynamic error; identical arguments apply to the dynamic distance and dynamic fidelity. The account is based upon [131].

The scenario we wish to consider is that of a system which is part of another, possibly much larger, system. For example, we may be interested in the performance of a single qubit memory element in a large quantum computer. One could argue that what should be done is to look at the fidelity of the total system – qubit plus computer. However, in general, quantum computers can be very large systems compared to the subsystem whose performance as a memory element we wish to analyze, and inclusion of the entire state and dynamics of the quantum computer would make the analysis enormously complicated.

Given that we do not wish to analyze the complete dynamics of the total system, the natural thing to do is to define a quantity which captures the worst-case error possible in the system. We define

$$E_1(\rho, \mathcal{E}) \equiv \max_{\tilde{\rho}, \mathcal{E}'} E\Big((\mathcal{E}' \otimes \mathcal{I}_Q)(\tilde{\rho}), (\mathcal{E}' \otimes \mathcal{E})(\tilde{\rho})\Big), \tag{5.113}$$

where the maximization is over all extensions $\tilde{\rho}$ of ρ to larger systems RQ, and all possible complete quantum operations $\mathcal{E}'$ that could occur on R. E_1 is a measure of how well the subsystem plus its

entanglement with the remainder of the system is stored. Note especially that the initial state RQ is not necessarily a pure state; it can be any extension of ρ whatsoever.

We maximize over all possible extensions and dynamics for the remainder of the system in order to obtain the *worst possible value* the error could have, regardless of the actual state or dynamics of the remainder of the system, R. The advantage of this is that this quantity depends only on the part of the computer, Q, under consideration, not on the detailed dynamics and state of the entire computer.

A second, related, quantity is also a useful measure of how well a system plus entanglement is stored. It will turn out that this quantity is equal to E_1. Define

$$E_2(\rho, \mathcal{E}) := \max_{\tilde{\rho}} E\Big(\tilde{\rho}, (\mathcal{I}_R \otimes \mathcal{E})(\tilde{\rho})\Big). \tag{5.114}$$

The motivation for this quantity is similar to that for E_1, except now we assume that R is subject to the identity dynamics $\mathcal{I}$, instead of maximizing over all possible dynamics $\mathcal{E}'$ for R.

To see that E_1 and E_2 are equal, note that

$$E_1(\rho, \mathcal{E}) \geq E_2(\rho, \mathcal{E}), \tag{5.115}$$

since the maximization in E_1 clearly includes all the values being maximized over for E_2. To see the reverse inequality, notice that

$$E\Big(\tilde{\rho}, (\mathcal{I}_R \otimes \mathcal{E})(\tilde{\rho})\Big) \geq E\Big((\mathcal{E}' \otimes \mathcal{I}_Q)(\tilde{\rho}), (\mathcal{E}' \otimes \mathcal{E})(\tilde{\rho})\Big), \tag{5.116}$$

by the non-increasing property of the error under quantum operations, and thus

$$E_2(\rho, \mathcal{E}) \leq E_1(\rho, \mathcal{E}). \tag{5.117}$$

It follows that

$$E_1(\rho, \mathcal{E}) = E_2(\rho, \mathcal{E}). \tag{5.118}$$

A similar argument can be used to show that $E(\rho, \mathcal{E}) = E_2(\rho, \mathcal{E})$. First, note that $E_2(\rho, \mathcal{E}) \geq E(\rho, \mathcal{E})$, by choosing the initial extension RQ to be a purification of Q. Second, $E_2(\rho, \mathcal{E}) \leq E(\rho, \mathcal{E})$, by the non-increasing property of the error under partial traces, which completes the proof.

It follows that

$$E(\rho, \mathcal{E}) = \max_{\rho, \mathcal{E}'} E\left((\mathcal{E}' \otimes \mathcal{I}_Q)(\tilde{\rho}), (\mathcal{E}' \otimes \mathcal{E})(\tilde{\rho})\right). \tag{5.119}$$

This expression brings home the operational meaning of the dynamic error in a way that is, perhaps, somewhat more compelling than the original abstract definition in terms of purifications, because it emphasizes the dynamic error as a quantity which arises as a *worst-case* scenario in contexts where preservation of entanglement may be important.

Summary of Chapter 5: Distance measures for quantum information

- **Absolute distance:** $D(\rho, \sigma) \equiv \mathrm{tr}|\rho - \sigma|$. Doubly convex metric on density operators, contractive under quantum operations.
- **Fidelity:**
$$F(\rho, \sigma) \equiv \mathrm{tr}\sqrt{\rho^{1/2}\sigma\rho^{1/2}} = \max_{|\psi\rangle, |\phi\rangle} |\langle\psi|\phi\rangle|.$$
Strongly concave, $F(\sum_i p_i\rho_i, \sum_i q_i\sigma_i) \geq \sum_i \sqrt{p_i q_i} F(\rho_i, \sigma_i)$.
- **Dynamic fidelity and dynamic distance:** $F(\rho, \mathcal{E})$ and $D(\rho, \mathcal{E})$. Measure how well entanglement is preserved during a quantum mechanical process, starting with the state ρ of a system Q, which is assumed to be entangled with another quantum system, R, and applying the quantum operation $\mathcal{E}$ to system Q.
- **Chaining of errors:** $E(\rho, \mathcal{E}_2 \circ \mathcal{E}_1) \leq E(\rho, \mathcal{E}_1) + E(\rho', \mathcal{E}_2)$, and similarly for the dynamic distance.

Part II

Bounds on quantum information transmission

Chapter 6

Quantum communication complexity

Suppose a number of widely separated parties wish to perform a distributed computation. Each of the parties has access to some part of the data which is to be used as input to the computation. However, no party has access to all of the data, so in general no party can complete the computation on their own. The *communication complexity* of a problem is defined to be the minimal communication cost incurred in performing the distributed computation. The classical theory of communication complexity was initiated by Yao [197], and has since blossomed into a dynamic field of research, as may be seen by consulting one of the excellent surveys of the field that have been written, such as [104]. Recently, Yao [198] has initiated the study of *quantum communication complexity*, in which quantum resources may be used to assist in the performance of a distributed computation.

The purpose of this Chapter is to develop some elementary results in quantum communication complexity. We will explore several different models for quantum communication complexity, in which different types of quantum resources may be used for communication, and with different computational goals. The models may be divided into two broad classes. The first class is concerned with the *quantum communication complexity of classical functions*. The problems in this class concern the computation of classical functions, but with quantum resources allowed to assist in the computation. We will examine several variants of this class, differentiated by the nature of the quantum resources used. For example, Yao [198] considered the computation of a classical function assisted by the ability of the computing parties to communicate using qubits. An important variant of this model was introduced by Cleve and Buhrman [44], who considered the computation of a classical function in which the communication is carried out using classical bits, but in which an arbitrary pre-shared entanglement is allowed. Other variants within this class will be mentioned during this Chapter.

The second class of models concern *coherent quantum communication complexity*. In this class, the problems involve the distributed computation of a *quantum function*, such as a joint unitary operation performed by Alice and Bob on their qubits. To my knowledge, this class of problems has not been discussed prior to this Dissertation.

The structure of the Chapter is as follows. Section 6.1 reviews the work of other researchers on the Holevo bound, an important result in quantum information theory, and the keystone of much of the later work in this chapter. Section 6.2 presents a complete, original solution to a basic problem in quantum information theory: what quantum resources are required to transmit classical information from one location to another, in the absence of noise? This result is used in section

6.3 to give an original lower bound on the quantum communication complexity of an interesting distributed computation, known as the *inner product problem*. The work reported in sections 6.2 and 6.3 is the result of a collaboration with Cleve, van Dam, and Tapp [47]. Section 6.4 reports the first results on coherent quantum communication complexity. I demonstrate a lower bound on the coherent quantum communication complexity of an important unitary operation, the quantum Fourier transform. Furthermore, a new and seemingly quite powerful general technique for proving lower bounds to the coherent quantum communication complexity is proved, and applied to several problems in coherent quantum communication complexity. Some of the results in subsection 6.4.2 were inspired by a conversation with Manny Knill. Section 6.5 outlines an original formalism which can be used to unify coherent quantum communication complexity with the quantum communication complexity for computing a classical function. The Chapter concludes in section 6.6 with a survey of some open problems in quantum communication complexity, and suggestions for future research directions.

My especial thanks to Richard Cleve for the many enjoyable and stimulating discussions about quantum communication complexity which stirred my interest in the field, and provoked many of the thoughts reported in this Chapter.

6.1 The Holevo bound

We begin with a review of what is historically perhaps the first major result in quantum information theory, the *Holevo bound* [83]. This result will be the basis for our later results in quantum communication complexity. The line of proof used here follows Schumacher, Westmoreland, and Wootters [159].

The setting is a game to be played by two fictitious protagonists, Alice and Bob. Alice is in possession of a classical source producing symbols $X = 1, \ldots, n$ according to a corresponding probability distribution $p_1, \ldots, p_n$. The aim of the game is for Alice to convey the value of X to Bob. However, for some reason, Alice can't give X directly to Bob. Rather, she prepares the quantum state ρ_X, where ρ_X is chosen from some fixed set $\rho_1, \ldots, \rho_n$ of quantum states. She then gives that state to Bob, whose task it is to determine the value of X, as best he can.

Suppose Bob performs a measurement on the quantum system he has been given, with measurement result Y. A measure of how much information he has gained about X is the mutual information $H(X:Y)$ discussed in Chapter 4. By the data processing inequality we know that Bob can infer X from Y if and only if $H(X:Y) = H(X)$, and that in general $H(X:Y) \leq H(X)$. More generally, it is true that the closer $H(X:Y)$ is to $H(X)$, the better Bob can do at inferring X from the observed value of Y. Bob's goal, therefore, is to choose a measurement which maximizes $H(X:Y)$, bringing it as close as possible to $H(X)$.

The Holevo bound states that:

$$H(X:Y) \leq S(\rho) - \sum_x p_x S(\rho_x), \tag{6.1}$$

where $\rho \equiv \sum_x p_x \rho_x$. Thus, the Holevo bound is an upper bound on the mutual information between Alice's classical data, X, and the result of Bob's measurement, Y. This bound holds *for any* measurement Bob may choose to do.

Before we proceed to the details of the proof, it is useful to note a few elementary formulas concerning the probabilities of various events. Suppose Bob does a measurement whose statistics are described by POVM elements M_y, corresponding to the different possible values which Y may

take. Then the probability that $Y = y$, given that the state ρ_x was prepared is given by

$$p(Y = y|X = x) = \mathrm{tr}(M_y \rho_x). \tag{6.2}$$

Thus $p(X = x, Y = y) = \mathrm{tr}(M_y \rho_x) p_x$, from which we can calculate $H(X, Y), H(X), H(Y)$ and thus $H(X : Y)$.

Our proof of Holevo's bound is not quite the most direct possible, however the route we take allows us to prove several facts that will be useful later in the Chapter. Our proof of Holevo's bound makes use of the following result:

Theorem 17 (Partial trace property of χ) *[159]*

Suppose states ρ_x of a system A are prepared, with respective probabilities p_x. Define the Holevo χ quantity for system A,

$$\chi_A \equiv S(\rho) - \sum_x p_x S(\rho_x), \tag{6.3}$$

where $\rho \equiv \sum_x p_x \rho_x$. Suppose a quantum system consists of two parts, A and B, and $\{p_x, \rho_x\}$ is an ensemble of states for the joint system AB. Then

$$\chi_A \leq \chi_{AB}, \tag{6.4}$$

where χ_{AB} and χ_A are the natural χ quantities associated with the ensemble $\{p_x, \rho_x\}$ for systems AB and A, respectively.

Proof [159]

Introduce a system, P, with an orthonormal basis $|x\rangle$ of states with index x corresponding to the index of the states ρ_x. Suppose the initial state of PAB is

$$\rho^{PAB} \equiv \sum_x p_x |x\rangle\langle x| \otimes \rho_x^{AB}. \tag{6.5}$$

Applying the joint entropy theorem on page 75 and doing a little algebra, we see that $\chi_{AB} = -S(P : A, B)$ and $\chi_A = -S(P : A)$. The result now follows from the observation that $S(P : A, B) \leq S(P : A)$, which is a restatement of strong subadditivity.

QED

It is straightforward to generalize this result to the case where a complete quantum operation, $\mathcal{E}$, replaces the partial trace in the above theorem. Suppose the system of interest is labeled Q, and let ρ_x'', ρ'' be the states obtained from ρ_x, ρ by applying $\mathcal{E}$, and let χ and χ'' be the corresponding Holevo χ quantities before and after application of the quantum operation $\mathcal{E}$. It is a simple corollary of the previous result that $\chi'' \leq \chi$. To see this, introduce a model environment E for the quantum operation $\mathcal{E}$. Suppose U^{QE} is the model interaction giving rise to the operation $\mathcal{E}$, and $|0\rangle$ is the initial state of E. Let ρ_x' be the *joint states* of QE after the model unitary operator has been applied, with ρ' and χ' defined in the obvious way. From the unitary invariance of the entropy, $\chi' = \chi$. But the states ρ_x'' may be obtained from the states ρ_x' by tracing out E, so by the previous theorem, $\chi'' \leq \chi$, as we set out to prove. We state this as a theorem generalizing the previous theorem:

Theorem 18 (Non-increasing property of χ under complete quantum operations) *[159]*

The Holevo χ quantity can not be increased under complete quantum operations.

The proof of the Holevo bound is to combine the partial trace property of χ with a beautiful construction involving four quantum systems, which we shall label P, Q, M and E. The interesting thing about the proof is that none of these systems need be associated with the "reality" of the problem at all; that is, these systems need not be directly related to the preparer, quantum system or observer appearing in the statement of Holevo's theorem; recall the use of similar constructions to prove entropic results in Chapter 4. The reason we can do this is because Holevo's theorem is an inequality between entropic quantities which do not depend on particular realizations for their meaning.

P is to be thought of as the "preparation" system. It has an orthonormal basis $|x\rangle$ whose elements correspond to possible preparations ρ_x for the quantum system, Q. M and E start out in standard pure states, which we will label $|0\rangle$ for both systems. The initial state of the total system is assumed to be

$$\rho^{PQME} = \sum_x p_x |x\rangle\langle x| \otimes \rho_x \otimes |0\rangle\langle 0| \otimes |0\rangle\langle 0|. \tag{6.6}$$

The intuition behind this construction is that system P represents Alice, the preparer, who knows the value of x, and depending on this value prepares an appropriate state for system Q. System M represents Bob's measuring apparatus, which records the result of the measurement, and E represents an additional "environmental record" [206] of this measurement. Formally, this measurement process is represented by a unitary dynamics on the system $PQME$ defined by the equation

$$U|PQ\rangle|0\rangle|0\rangle \equiv \sum_y \left(I^P \otimes \sqrt{M_y}\right) |PQ\rangle|y\rangle|y\rangle, \tag{6.7}$$

where $|PQ\rangle$ is any pure state of PQ. As we saw in Chapter 3 the unitarity of the dynamics defined by this equation is ensured by the completeness relation $\sum_y M_y = I$. The state of the system after this evolution is

$$\rho^{PQME'} = \sum_{x,y_1,y_2} p_x |x\rangle\langle x| \otimes \left(\sqrt{M_{y_1}}\rho_x\sqrt{M_{y_2}}\right) \otimes |y_1\rangle\langle y_2| \otimes |y_1\rangle\langle y_2|. \tag{6.8}$$

Let χ'_M be the Holevo χ quantity associated with system M after the interaction. The respective states of system M after the interaction, are given by

$$M'_x = \sum_y p(y|x)|y\rangle\langle y| \tag{6.9}$$

$$M' = \sum_y p(y)|y\rangle\langle y|, \tag{6.10}$$

so $\chi'_M = H(Y) - H(Y|X) = H(X:Y)$. By the partial trace property, $H(X:Y) = \chi'_M \leq \chi'_{QME}$. But the interaction of Q, M and E was unitary, so $\chi'_{QME} = \chi_{QME} = \chi_Q$. Putting it all together, we see that

$$H(X:Y) \leq \chi_Q = S(\rho) - \sum_x p_x S(\rho_x), \tag{6.11}$$

which is the Holevo bound.

Holevo's bound is a keystone in the proof of many results in quantum information theory. The remainder of this section samples some of the well-known uses to which the Holevo bound may be put, in order to sharpen your intuition about the uses of this result. Recall from page 77 that

$$S(\rho) \leq H(X) + \sum_x p_x S(\rho_x), \tag{6.12}$$

with equality if and only if the states ρ_x have orthogonal support. Suppose that the states ρ_x do not have orthogonal support. Then Holevo's theorem allows us to conclude that $H(X:Y)$ is always strictly less than $H(X)$. This is just our intuitive notion that if the states prepared by Alice are not orthogonal, then it is not possible for Bob to determine with certainty which state Alice prepared.

A more concrete example may be useful. Suppose Alice prepares a single qubit in one of two quantum states. Which state she prepares is determined by a fair coin toss. If the coin toss yields heads, then Alice prepares the state $|0\rangle$, and if the coin toss yields tails, then Alice prepares the state $\cos\theta|0\rangle + \sin\theta|1\rangle$, where θ is some real parameter. In the $|0\rangle, |1\rangle$ basis it follows that ρ can be written

$$\rho = \frac{1}{2}\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} + \frac{1}{2}\begin{bmatrix} \cos^2\theta & \cos\theta\sin\theta \\ \cos\theta\sin\theta & \sin^2\theta \end{bmatrix}. \tag{6.13}$$

A simple calculation shows that the eigenvalues of ρ are $(1 \pm \cos\theta)/2$. This allows us to calculate the Holevo bound, which in this case of pure state signals, is just the entropy of ρ.

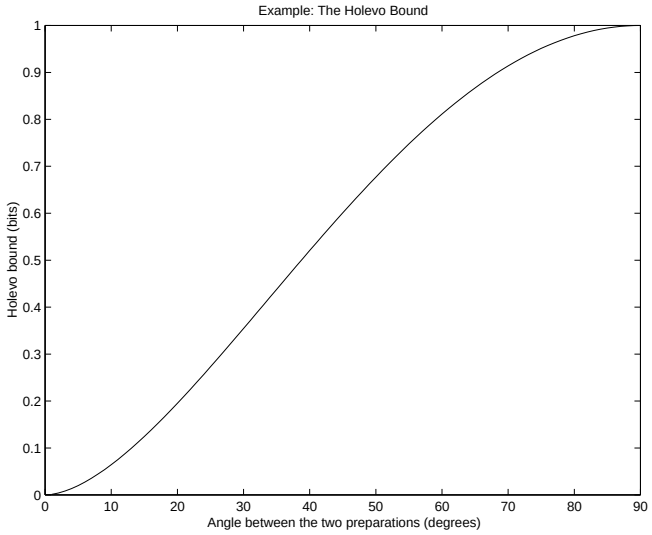

Figure 6.1: Plot of the Holevo bound when the states $|0\rangle$ and $\cos\theta|0\rangle + \sin\theta|1\rangle$ are prepared with equal probability. Notice that the Holevo bound reaches a maximum when the angle between the two states is $\theta = \pi/2$, corresponding to orthogonal states. At this point only it is possible for Bob to determine with certainty which state Alice prepared.

A plot of the Holevo bound for this example is shown in figure 6.1. Notice that the Holevo bound is maximized at one bit, when θ is 90 degrees, corresponding to orthogonal states. At this point it is possible for Bob to determine with surety which state Alice prepared. For other values of the θ, the Holevo bound is strictly less than one bit, and it is impossible for Bob to determine with surety which state Alice prepared, since $H(X)$, Alice's preparation entropy, is equal to one bit.

This can be quantified more precisely by making use of the *Fano inequality*. The Fano inequality is a result of classical information theory, proved in the box on page 115, which provides a connection between the loss of mutual information and the likelihood that an error is made in

inference. Suppose Bob makes a guess $\tilde{X} = f(Y)$ as to which state Alice prepared, based on the outcome of his measurement, Y. In the general case, according to the Fano inequality,

$$H_{\rm bin}(p(\tilde{X} \neq X)) + p(\tilde{X} \neq X)) \log(|X| - 1) \geq H(X|Y) = H(X) - H(X:Y). \tag{6.14}$$

Combining this result with the Holevo bound allows us to place bounds on how well Bob may infer the value of X. Heuristically, the smaller χ is, the harder it is for Bob to determine which state Alice prepared. More precisely, we have

$$S(\rho) - \sum_x p_x S(\rho_x) + H_{\rm bin}(p(\tilde{X} \neq X)) + p(\tilde{X} \neq X)) \log(|X| - 1) \geq H(X). \tag{6.15}$$

We can use this equation to numerically set a lower bound on the probability of Bob making an error in inference of Alice's original state. For instance, in the present example, where Alice prepares $|0\rangle$ with probability one half and $\cos\theta|0\rangle + \sin\theta|1\rangle$ with probability one half, the inequality reduces to

$$H_{\rm bin}(p(\tilde{X} \neq X)) \geq 1 - H\left(\frac{1+\cos\theta}{2}\right). \tag{6.16}$$

A little thought shows that when $\theta \neq 90$ degrees, we must have $p(\tilde{X} \neq X) > 0$; moreover, the further away θ is from 90 degrees, the larger the probability of an error in inference is constrained to be by this inequality.

Box 6.1: Fano's inequality

Suppose we wish to infer the value of an unknown random variable, X, based on knowledge of another random variable, Y. Intuitively, we expect that the conditional entropy $H(X|Y)$ limits how well we may perform this inference. The *Fano inequality* [54] provides a useful bound on how well we may infer X, given Y.

Suppose $\tilde{X} = f(Y)$ is some function of Y which we are using as our best guess for X. Let $p_e \equiv p(\tilde{X} \neq X)$ be the probability that this guess is incorrect. Then the Fano inequality states that

$$H_{\rm bin}(p_e) + p_e \log(|X| - 1) \geq H(X|Y), \tag{6.17}$$

where $H_{\rm bin}$ is the binary entropy and $|X|$ is the number of values X may assume. Examining the inequality, what it tells us is that if $H(X|Y)$ is large, then the probability of making an error in inference, p_e, must also be large.

To prove the Fano inequality, define an "error" random variable, $E \equiv 1$ if $X \neq \tilde{X}$, and $E \equiv 0$ if $X = \tilde{X}$. Notice that $H(E) = H_{\rm bin}(p_e)$. Using the chain rule for entropies proved on page 72, we have $H(E, X|Y) = H(E|X,Y) + H(X|Y)$. But E is completely determined once X and Y are known, so $H(E|X,Y) = 0$ and thus $H(E,X|Y) = H(X|Y)$. Applying the chain rule for entropies in a different fashion, we obtain $H(E,X|Y) = H(X|E,Y) + H(E|Y)$. Conditioning reduces entropy, so $H(E|Y) \leq H(E) = H_{\rm bin}(p_e)$. Finally,

$$\begin{aligned} H(X|E,Y) &= p(E=0)H(X|E=0,Y) + p(E=1)H(X|E=1,Y) & (6.18)\\ &\leq p(E=0) \times 0 + p_e \log(|X| - 1), & (6.19)\end{aligned}$$

where $H(X|E=1,Y) \leq \log(|X|-1)$ follows from the fact that when $E=1$, $X \neq Y$, and X can assume at most $|X|-1$ values, bounding its entropy, and thus its conditional entropy by $\log(|X|-1)$. Summarizing, we have

$$H(X|Y) = H(E,X|Y) \leq H_{\rm bin}(p_e) + p_e \log(|X| - 1), \tag{6.20}$$

which establishes the Fano inequality.

6.2 Capacity theorem for qubit communication

We can use the Holevo bound to analyze the following two-party game. Alice is in possession of n bits which she would like to transmit to Bob. To achieve this, she is allowed to send qubits to Bob, and Bob may send qubits to Alice, with no other form of communication allowed. How many qubits must Alice and Bob use in order to successfully transmit the n bits from Alice to Bob? The following *capacity theorem* provides a complete answer to this question:

Theorem 19 (Capacity theorem for communication using qubits)
Suppose that Alice possesses n bits of information, and wants to convey this information to Bob. Suppose that Alice and Bob possess no prior entanglement but qubit communication in either direction is allowed. Let n_{AB} be the number of qubits Alice sends to Bob, and n_{BA} the number of qubits Bob sends to Alice. Then, Bob can acquire the n bits if and only if the following inequalities are satisfied:

$$n_{AB}, n_{BA} \geq 0 \tag{6.21}$$
$$n_{AB} \geq \lceil n/2 \rceil \tag{6.22}$$
$$n_{AB} + n_{BA} \geq n. \tag{6.23}$$

Moreover, the necessity of the condition $n_{AB} \geq \lceil n/2 \rceil$ remains valid even if pre-shared entanglement is allowed. More generally, Bob can acquire m bits of mutual information with respect to Alice's n bits if and only if the above equations hold with m substituted for n.

Graphically, the capacity region for the above communication problem is shown in figure 6.2. Note the difference with the classical result for communication with bits, where the capacity region is given by the equation $n_{AB} \geq n$; that is, classically, communication from Bob to Alice does not help.

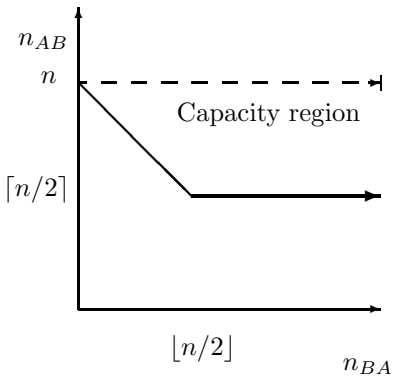

Figure 6.2: Capacity region to send n bits from Alice to Bob. n_{AB} is the number of qubits Alice sends to Bob, and n_{BA} is the number of qubits Bob sends to Alice. The dashed line indicates the bottom of the classical capacity region.

Proof

The sufficiency of equations (6.22) and (6.23) follows from the superdense coding technique discussed in section 2.2. Sufficiency in the case $n_{AB} \geq n$ is obvious, so we suppose $n_{AB} < n$. Bob prepares $n - n_{AB} \leq n_{BA}$ maximally entangled pairs of qubits, and sends one qubit of each

pair to Alice, who can use them in conjunction with sending $n - n_{AB} \leq n_{AB}$ qubits to Bob to transmit $2(n - n_{AB})$ bits to Bob, using superdense coding. Alice uses her remaining allotment of $n_{AB} - (n - n_{AB}) = 2n_{AB} - n \geq 0$ qubits to transmit the remaining $2n_{AB} - n$ bits in the obvious way.

The proof that equations (6.22) and (6.23) are necessary follows from an application of Holevo's theorem and the non-increasing property of the Holevo χ under partial traces, as discussed in the previous section. The details are as follows.

Let X be Alice's n bits of information, which we assume is uniformly distributed over $\{0,1\}^n$. Without loss of generality, it can be assumed that the protocol between Alice and Bob is of the following form. For any value $(x_1, \ldots, x_n)$ of X, Alice begins with a set of qubits in state $|x_1, \ldots, x_n\rangle|0, \ldots, 0\rangle$ and Bob begins with a set of qubits in a standard state $|0, \ldots, 0\rangle$. The protocol consists of a sequence of steps, where at each step one of the following four processes takes place.

1. Alice performs a unitary operation on the qubits in her possession.
2. Bob performs a unitary operation on the qubits in his possession.
3. Alice sends a qubit to Bob.
4. Bob sends a qubit to Alice.

After these steps, Bob performs a measurement on the qubits in his possession, which has outcome Y. Bob's goal is to maximize the mutual information between Y and Alice's input, X. You may be wondering about the possibility of a protocol in which Bob starts with a mixed state, or in which non-unitary operations are allowed. Note that using the techniques of Chapter 3 any non-unitary operation can be simulated by a unitary operation, with the introduction of an appropriate environmental model, and by the purification procedure discussed in Appendix A it is possible to simulate any protocol in which Bob starts with a mixed state by a protocol in which Bob starts with a pure state.

Let ρ_i^X be the density operator of the set of qubits that are in Bob's possession after i steps in the protocol have been executed, and $\rho_i \equiv \sum_x p_x \rho_i^x$ be the density operator of Bob's system after i steps, averaged over all possible inputs, x. Due to Holevo's bound, it suffices to upper bound the final value of $\chi(\rho_i^X)$. We consider the evolution of $\chi(\rho_i^X)$ and $S(\rho_i)$. Initially, $\chi(\rho_0^X) = S(\rho_0) = 0$, since Bob begins in a state independent of X. Now, consider how $\chi(\rho_i^X)$ and $S(\rho_i)$ change for each of the four processes above.

1. Alice performing a unitary operation on her qubits does not affect ρ_i^X and hence has no effect on $\chi(\rho_i^X)$ or $S(\rho_i)$.
2. It is easy to verify that χ and S are invariant under unitary transformations, so Bob performing a unitary on his qubits does not affect $\chi(\rho_i^X)$ and $S(\rho_i)$, either.
3. Alice sends a qubit to Bob. Let B denote Bob's qubits after i steps and Q denote the qubit that Alice sends to Bob at step number $i+1$. By the subadditivity inequality discussed in subsection 4.3.5 and the fact that, for a single qubit Q, $S(Q) \leq 1$, $S(B,Q) \leq S(B) + S(Q) \leq S(B) + 1$. Also, by the triangle inequality discussed in subsection 4.3.5, $S(B,Q) \geq S(B) - S(Q) \geq S(B) - 1$. It follows that $S(\rho_{i+1}) \leq S(\rho_i) + 1$ and thus

$$\begin{aligned} \chi(\rho_{i+1}^X) &= S(\rho_{i+1}) - \sum_x p_x S(\rho_{i+1}^x) \\ &\leq (S(\rho_i) + 1) - \sum_x p_x (S(\rho_x) - 1) \end{aligned}$$

$$= \quad \chi(\rho_i^X) + 2. \tag{6.24}$$

4. Bob sends a qubit to Alice. In this case, ρ_{i+1}^X is ρ_i^X with one qubit traced out. We saw in the previous section that χ does not increase under partial trace, so $\chi(\rho_{i+1}^X) \leq \chi(\rho_i^X)$. Note also that $S(\rho_{i+1}) \leq S(\rho_i) + 1$ for this process, by the triangle inequality.

Now, since $\chi(\rho_i^X)$ can only increase when Alice sends a qubit to Bob and by at most 2, equation (6.22) follows from the Holevo bound. Also, since $S(\rho_i)$ can only increase when one party sends a qubit to the other and by at most 1, equation (6.23) follows from the observation that $S(\rho_i)$ is an upper bound on the Holevo χ, and the Holevo bound. Finally, note that even if pre-shared entanglement is allowed, so $S(\rho_i)$ may start out greater than zero, $\chi(\rho_i)$ is still zero at the start of the protocol, and thus the reasoning leading to the constraint $n_{AB} \geq \lceil n/2 \rceil$ still holds. This completes the proof.

QED

Our main interest in this capacity theorem is as a step along the way to proving results about quantum communication complexity. However, it is interesting to briefly consider other directions in which this work could be taken.

First, let us consider what the essential difference is between the classical and the quantum resources required to perform the task under consideration. Recall that when Alice sends a qubit to Bob, we showed that $\Delta\chi \leq 2$. To prove this, we used the triangle inequality $S(B,Q) \geq S(B) - S(Q)$ to show that $S(\rho'_x) \geq S(\rho_x) - 1$. However, in the absence of entanglement between Q and B, we will show in Chapter 8 that the stronger inequality $S(B,Q) \geq S(B)$ is true. We deduce that in the absence of entanglement, $\Delta\chi \leq 1$, from which the familiar classical lower bound $n_{AB} \geq n$ emerges.

Second, note that we have assumed noiseless transmission of quantum information between Alice and Bob. What are the resource requirements if there is noise in the channel between Alice and Bob? Another interesting path for generalization is to consider the many-party version of the problem. What quantum resources are required to accomplish communication of classical information amongst a network of k users? Finally, we may ask for a precise characterization of what quantum resources are required to transmit n bits of information in the presence of a pre-shared entanglement between Alice and Bob. Answering this question in full generality may give new insight into the meaning of entanglement, and suggests a means for defining measures of entanglement for quantum systems consisting of two or more components.

6.3 Communication complexity of the inner product

We now have the tools we need to investigate several interesting problems in quantum communication complexity. We begin with the communication complexity of the *inner product modulo two (IP)* function:

$$IP(x,y) \quad = \quad (x_1 \cdot y_1 + x_2 \cdot y_2 + \cdots + x_n \cdot y_n) (\mathrm{mod} 2). \tag{6.25}$$

The communication complexity of the IP function is fairly well understood in the classical models of communication complexity. For worst-case inputs and deterministic protocols guaranteed to give the correct answer, the communication complexity is n. For randomized protocols (with either an independent or a shared random source), uniformly distributed or worst-case inputs, and with error probability $\frac{1}{2} - \delta$ required, the communication complexity is $n - O(\log(1/\delta))$ [37] (see also [104]).

Yao [198] has introduced a model of quantum communication complexity in which Alice and Bob start of with uncorrelated systems, and are allowed to communicate using qubits, in order

to compute some joint classical function. The cost, or communication complexity, in this model, is defined to be the minimum number of qubits which must be communicated in order to compute the classical function. Kremer [103], using a proof methodology which he attributes to Yao, demonstrated that in this model the communication complexity of IP is asymptotically linear in n, whenever the required correctness probability is $1 - \epsilon$ for a constant $0 \leq \epsilon < \frac{1}{2}$.

In this section, we consider the communication complexity of IP in two models different to that introduced by Yao: with prior entanglement and qubit communication; and with prior entanglement and classical bit communication. In both models, an arbitrary prior entanglement may be set up, at no cost to the protocol. As far as is presently known, the proof methodology of the lower bound in the qubit communication model without prior entanglement [103] does not carry over to either of these two models. Nevertheless, we show linear lower bounds in these models.

To state our lower bounds more precisely, we introduce the following notation. Let $f : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$ be a communication problem, and $0 \leq \epsilon < \frac{1}{2}$. Let $Q^*_\epsilon(f)$ denote the communication complexity of f in terms of *qubits*, where quantum entanglement is available and the requirement is that Bob determines the correct answer with probability at least $1 - \epsilon$; the $*$ superscript is intended to highlight the fact that prior entanglement is available. Also, let $C^*_\epsilon(f)$ denote the corresponding communication complexity of f in the scenario where the communication is in terms of *bits*; again, quantum entanglement is available and Bob is required to determine the correct answer with probability at least $1 - \epsilon$. When $\epsilon = 0$, we refer to the protocols as *exact*, and, when $\epsilon > 0$, we refer to them as *bounded-error* protocols. With this notation, the results we will prove in this section may be summarized:

$$Q^*_0(\mathrm{IP}) = \lceil n/2 \rceil \tag{6.26}$$

$$Q^*_\epsilon(\mathrm{IP}) \geq \frac{1}{2}(1-2\epsilon)^2 n - \frac{1}{2} \tag{6.27}$$

$$C^*_0(\mathrm{IP}) = n \tag{6.28}$$

$$C^*_\epsilon(\mathrm{IP}) \geq \max(\frac{1}{2}(1-2\epsilon)^2, (1-2\epsilon)^4) n - \frac{1}{2} \tag{6.29}$$

Note that all the lower bounds are linear in n whenever ϵ is held constant. Also, these results subsume the lower bounds in [103], since the qubit model defined by Yao [198] differs from the bounded-error qubit model defined above only in that it does not permit a prior entanglement.

The lower bound proofs employ a novel kind of "quantum" reduction between protocols, which reduces the problem of communicating, say, n bits of information to the IP problem. It is noteworthy that there does not appear to be a similar classical reduction between the two problems. This reduction is particularly remarkable since quantum information theory subsumes classical information theory, and therefore our results also represent new proofs of nontrivial lower bounds on the *classical* communication complexity of IP. It is intriguing that we are able to prove such lower bounds using a quantum mechanical methodology fundamentally different from previous methods used for proving classical lower bounds.

6.3.1 Converting exact protocols into clean form

We begin by showing how to reduce a general protocol for computing a function $f(x, y)$ into a special type of protocol which we call a *clean protocol*. A clean protocol is a special kind of qubit protocol inspired by the general spirit of the reversible computing paradigm [110, 14], in a quantum setting. In particular, a clean protocol is set up so that none of the qubits involved in the protocol changes, except for one, which contains the answer, $f(x, y)$.

In general, the initial state of a qubit protocol is of the form

$$\underbrace{|y_1,\ldots,y_n\rangle \otimes |0,\ldots,0\rangle}_{\text{Bob's qubits}} \otimes |\Phi_{BA}\rangle \underbrace{\otimes |x_1,\ldots,x_n\rangle \otimes |0,\ldots,0\rangle}_{\text{Alice's qubits}}, \tag{6.30}$$

where $|\Phi_{BA}\rangle$ is the state of the entangled qubits shared by Alice and Bob, and the $|0,\ldots,0\rangle$ states can be regarded as work space for the protocol. At each turn, one party performs some unitary operation on all the qubits in their possession and then sends a subset of these qubits to the other party. Note that, due to the communication, the set of qubits possessed by each party varies during the execution of the protocol.

We say that a protocol which exactly computes a function $f(x,y)$ is *clean* if, when executed on the initial state

$$|z\rangle \otimes |y_1,\ldots,y_n\rangle \otimes |0,\ldots,0\rangle \otimes |\Phi_{BA}\rangle \otimes |x_1,\ldots,x_n\rangle \otimes |0,\ldots,0\rangle, \tag{6.31}$$

the protocol results in the final state

$$|z+f(x,y)\rangle \otimes |y_1,\ldots,y_n\rangle \otimes |0,\ldots,0\rangle \otimes |\Phi_{BA}\rangle \otimes |x_1,\ldots,x_n\rangle \otimes |0,\ldots,0\rangle \tag{6.32}$$

(where the addition is mod 2). The input, the work qubits, and the initial entangled qubits will typically change states during the execution of the protocol, but they are reset to their initial values at the end of the protocol.

We will show how to transform a general protocol into a clean protocol. This transformation comes at a cost, which we quantity using the following notation. If a qubit protocol consists of m_1 qubits from Alice to Bob and m_2 qubits from Bob to Alice then we refer to the protocol as an (m_1,m_2)-qubit protocol.

The following argument shows it is always possible to transform an exact (m_1,m_2)-qubit protocol to compute $f(x,y)$ into a clean (m_1+m_2,m_1+m_2)-qubit protocol that computes the same function. First, the protocol for f is run once, creating a state of the form

$$|z\rangle|f(x,y)\rangle \otimes |\Phi'_{BA}\rangle\rangle, \tag{6.33}$$

where $|\Phi'_{BA}\rangle$ is some extra "garbage" state of the joint system, BA, which will depend on $x_1,\ldots,x_n,y_1,\ldots,y_n$. We now apply a controlled not gate with $|f(x,y)\rangle$ as the control and $|z\rangle$ as the ancilla, to create the state

$$|z+f(x,y)\rangle|f(x,y)\rangle \otimes |\Phi'_{BA}\rangle. \tag{6.34}$$

Finally, note that all the steps in the protocol to compute f were reversible. We now run the original protocol in reverse, putting the system in the desired state,

$$|z+f(x,y)\rangle \otimes |y_1,\ldots,y_n\rangle \otimes |0,\ldots,0\rangle \otimes |\Phi_{BA}\rangle \otimes |x_1,\ldots,x_n\rangle \otimes |0,\ldots,0\rangle. \tag{6.35}$$

Note that, for each qubit that Bob sends to Alice when the protocol is run forwards, Alice sends the qubit to Bob when run in the backwards direction. Therefore, we have constructed a (m_1+m_2,m_1+m_2)-qubit protocol that maps state (6.31) to state (6.32).

6.3.2 Reduction from the communication problem

We now show how to transform a clean (m_1+m_2,m_1+m_2)-qubit protocol that exactly computes IP for inputs of size n, to an (m_1+m_2,m_1+m_2)-qubit protocol that transmits n bits of information from Alice to Bob. This is accomplished in four stages:

1. Bob initializes his qubits as indicated in equation (6.31) with $z = 1$ and $y_1 = \cdots = y_n = 0$, while Alice prepares the state $|x_1, \ldots, x_n\rangle$.

2. Bob performs a Hadamard gate on each of his first $n+1$ qubits.

3. Alice and Bob execute the clean protocol for the inner product function.

4. Bob again performs a Hadamard gate on each of his first $n+1$ qubits.

Let $|B_i\rangle$ denote the state of Bob's first $n+1$ qubits after the i^{th} stage. Recalling the definition of the Hadamard gate from page 21, $H|0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and $H|1\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$, we see that

$$
\begin{aligned}
|B_1\rangle &= |1\rangle \otimes |0, \ldots, 0\rangle && (6.36)\\
|B_2\rangle &= \tfrac{1}{\sqrt{2^{n+1}}} \sum_{a,b_1,\ldots,b_n \in \{0,1\}} (-1)^a |a\rangle \otimes |b_1, \ldots, b_n\rangle && (6.37)\\
|B_3\rangle &= \tfrac{1}{\sqrt{2^{n+1}}} \sum_{a,b_1,\ldots,b_n \in \{0,1\}} (-1)^a |a + b_1 x_1 + \cdots + b_n x_n\rangle \otimes |b_1, \ldots, b_n\rangle \\
&= \tfrac{1}{\sqrt{2^{n+1}}} \sum_{c,b_1,\ldots,b_n \in \{0,1\}} (-1)^{c + b_1 \cdot x_1 + \cdots + b_n \cdot x_n} |c\rangle \otimes |b_1, \ldots, b_n\rangle && (6.38)\\
|B_4\rangle &= |1\rangle \otimes |x_1, \ldots, x_n\rangle, && (6.39)
\end{aligned}
$$

where, in equation (6.38), the substitution $c = a + b_1 x_1 + \cdots + b_n x_n$ has been made, and arithmetic over bits is taken modulo 2. The above transformation was inspired by [173]; see also [25].

Since the above protocol conveys n bits of information (namely, $x_1, \ldots, x_n$) from Alice to Bob, by the capacity theorem of section 6.2, we have $m_1 + m_2 \geq n/2$. Since this protocol can be constructed from an arbitrary exact (m_1, m_2)-qubit protocol for IP, this establishes the lower bound of equation (6.26). That this bound is achievable follows immediately from the superdense coding technique; Alice need merely send all n of her bits to Bob using $\lceil n/2 \rceil$ qubits and superdense coding. Bob can then calculate the inner product. This completes the proof of equation (6.26).

The approximate result (6.27) follows in a straightforward fashion by running essentially the same argument, and using the Fano inequality to bound the probability of error for the IP protocol. We will not go through the details here; they may be found in [47].

Note that, classically, the reduction used here to prove the communication complexity lower bound is not possible. For example, if a clean protocol for IP is executed in any classical context, it can never yield more than one bit of information to Bob, whereas, in this quantum context, it yields n bits of information to Bob.

6.3.3 Lower bounds for bit protocols

We now use the just-proved exact *quantum bit* communication complexity for IP to prove an exact *classical bit* communication complexity for IP, in the presence of pre-shared entanglement, equation (6.28).

Using quantum teleportation it is straightforward to simulate any m-qubit protocol by a $2m$-bit classical protocol, and appropriate pre-shared entanglement. Also, if the communication pattern in an m-bit protocol is such that an even number of bits is always sent during each party's turn then it can be simulated by an $m/2$-qubit protocol by superdense coding [23] (which also employs EPR pairs). However, this latter simulation technique cannot, in general, be applied directly, especially for protocols where the parties take turns sending single bits.

We can nevertheless obtain a slightly weaker simulation of bit protocols by qubit protocols for IP that is sufficient for our purposes. The result is that, given any m-bit protocol for IP_n (that is, IP instances of size n), one can construct an m-qubit protocol for IP_{2n}. This is accomplished by interleaving two executions of the bit protocol for IP_n to compute two independent instances of inner products of size n. We make two observations. First, by taking the sum (mod 2) of the two results, one obtains an inner product of size $2n$. Second, due to the interleaving, an even number of bits is sent at each turn, so that the above superdense coding technique can be applied, yielding a $(2m)/2 = m$-qubit protocol for IP_{2n}. Now, equation (6.26) implies $m \geq n$, which establishes the lower bound of equation (6.28). The achievability of this lower bound follows from the obvious protocol: Alice sends her classical data to Bob, establishing equality in equation (6.28).

To obtain the lower bound (6.29), suppose we apply the same proof technique as above to any m-bit protocol computing IP_n with probability $1-\epsilon$. We obtain an m-qubit protocol which computes IP_{2n} with probability $(1-\epsilon)^2 = 1-2\epsilon(1-\epsilon)$. Applying equation (6.27), with $2n$ replacing n and $2\epsilon(1-\epsilon)$ replacing ϵ, we find that $m \geq (1-2\epsilon)^4 n - 1/2$. For $\epsilon > 1/2 - \sqrt{2}/4 \approx 0.146$ a better bound is obtained by noting that $C^*_\epsilon \geq Q^*_\epsilon$ is always true, since quantum bits can always be used in the place of bits, and applying equation (6.27). This establishes equation (6.29).

6.4 Coherent quantum communication complexity

In the previous section we considered the distributed computation of a classical function using quantum resources. Analogous questions can be asked about the distributed computation of a quantum function using quantum resources, a field of investigation which we will call *coherent quantum communication complexity*, or more usually just *coherent communication complexity*.

In this section we develop some elementary lower bounds on the coherent communication complexity. The original inspiration for this investigation was the following problem, which we shall call FT, for *Fourier transform*. Suppose Alice is in possession of n qubits, Bob is in possession of n qubits, and they wish to perform the quantum Fourier transform [163, 48, 59]. How many qubits must be communicated between Alice and Bob if they are to achieve this goal?

We will prove a lower bound of n qubits for this problem, using a method inspired by that used to prove the IP lower bound. We then prove a much more general lower bound, which applies to any unitary operator. This general lower bound is then used to give an alternate proof that the quantum Fourier transform has a coherent communication complexity of at least n qubits. The section closes with some general remarks about further directions for exploration in the field of coherent quantum communication complexity.

6.4.1 Coherent communication complexity of the quantum Fourier transform

It is clear that the computation of FT – and, indeed, of any unitary transform – can be done using $2n$ qubits of communication: Alice sends her n qubits to Bob, who performs the quantum Fourier transform on all $2n$ qubits. Bob then sends the n qubits which initially belonged to Alice back to Alice, completing the quantum Fourier transform. We will show that this is essentially the best procedure that can be achieved, to within a constant factor. Our general strategy will be to use a technique similar to that used for the inner product function, transforming a protocol for computing the quantum Fourier transform into a communications protocol.

It has been shown by Danielson and Lanczos (see [145] for a discussion) that the Fourier

transform on m qubits has the following effect:

$$|x_1, \ldots, x_m\rangle \longrightarrow \left(|0\rangle + e^{2\pi i x_m}|1\rangle\right) \otimes \left(|0\rangle + e^{2\pi i 0.x_{m-1}x_m}|1\rangle\right) \otimes \ldots \otimes \left(|0\rangle + e^{2\pi i 0.x_1 \ldots x_m}|1\rangle\right), \quad (6.40)$$

where $x_1, \ldots, x_m$ is any set of m bits. Throughout this section normalization factors are omitted. This decomposition, discovered by Danielson and Lanczos in 1942, has been rediscovered many times since; in the quantum context it has been rediscovered by Griffiths and Niu [75] and, somewhat later, but independently, by Cleve *et al* [46].

The strategy we use is to turn a protocol for computing the quantum Fourier transform into a method for classical communication between Alice and Bob. Suppose Alice has a string $x_1, \ldots, x_n$ of classical bits which she wishes to transmit to Bob. The following protocol achieves this. Alice prepares a system of n qubits in the state

$$|x_1, \ldots, x_n\rangle, \quad (6.41)$$

while Bob prepares a system of n qubits in the all $|0\rangle$ state.

Alice and Bob now jointly apply the $2n$-qubit quantum Fourier transform to their system, resulting in the state

$$(|0\rangle + |1\rangle)^{\otimes n} \otimes \left(|0\rangle + e^{2\pi i 0.x_n}|1\rangle\right) \otimes \left(|0\rangle + e^{2\pi i 0.x_{n-1}x_n}|1\rangle\right) \otimes \ldots \otimes \left(|0\rangle + e^{2\pi i 0.x_1 \ldots x_n}|1\rangle\right). \quad (6.42)$$

Bob now performs an n-qubit inverse quantum Fourier transform on his n qubits, resulting in a final state for Bob of $|x_1, \ldots, x_n\rangle$, from which he can simply read off the values of $x_1, \ldots, x_n$ which were originally in Alice's possession.

Thus any procedure for performing the quantum Fourier transform immediately yields a procedure for communicating n classical bits of information from Alice to Bob, for the same cost. The results of section 6.2 imply a lower bound on the coherent communication complexity of the quantum Fourier transform of n qubits. We conclude that the coherent communication complexity of the quantum Fourier transform is in the range n to $2n$ qubits.

So far we have considered the coherent communication complexity of the quantum Fourier transform in the case where the quantum Fourier transform must be done exactly. What if we are willing to allow an *approximate* performance of the quantum Fourier transform? Suppose we are attempting to perform the unitary operation U, but the protocol instead performs a quantum operation $\mathcal{E}$. Define the *absolute distance* for the protocol by

$$D \equiv D(U, \mathcal{E}) \equiv \min_{|\psi\rangle} D(U|\psi\rangle, \mathcal{E}(|\psi\rangle\langle\psi|)), \quad (6.43)$$

where the minimization is over all pure states $|\psi\rangle$, and the function $D(\cdot, \cdot)$ appearing on the right hand side is the absolute distance of Chapter 5. The absolute distance for the protocol is a measure of how well the protocol computes the quantum Fourier transform.

Suppose we want a protocol such that $D < \epsilon$. Intuitively, any such protocol must involve nearly n qubits, with the allowed deviation determined by the magnitude of ϵ.

Suppose q qubits are sent during the protocol for the approximate quantum Fourier transform. Suppose we substitute this approximate quantum Fourier transform into the bit communication protocol used earlier to obtain the exact lower bound. Since only q qubits are sent, we have an upper bound on the final Holevo χ quantity of Bob's system of q bits. Combining the Fano

inequality and the Holevo bound, this implies that the minimal probability of Bob making an error in his inference, p_e, satisfies

$$(n - q) \leq h(p_e) + p_e n. \tag{6.44}$$

But from page 91 we know that $p_e \leq D$, and assuming that $\epsilon < 1/2$, it follows that $h(p_e) \leq h(D)$, so

$$q \geq n(1 - \epsilon) - h(\epsilon). \tag{6.45}$$

In the case when $\epsilon = 0$ this reduces to $q \geq n$, as expected, and more generally gives a lower bound on the number of qubits required to achieve a specified accuracy in the coherent communication complexity.

It is somewhat disappointing that there remains the gap between n and $2n$ for the coherent communication complexity of the quantum Fourier transform. I have not been able to close this gap. I conjecture that $2n$ is the actual coherent communication complexity of the quantum Fourier transform; the proof or refutation of this conjecture is an interesting problem for future work.

6.4.2 A general lower bound

Let U be a general unitary operator on a joint system, AB. Suppose the action on A is on m qubits, and on B is on n qubits. Note that U can always be written in the form

$$U = \sum_i A_i \otimes B_i, \tag{6.46}$$

where A_i and B_i are non-zero operators on the systems A and B, respectively. The *Schmidt number* of U, $\mathrm{Sch}(U)$, is defined to be the minimal number of operators A_i (equivalently B_i) required in any such decomposition of U.

A general lower bound on the coherent communication complexity of U is

$$Q_0(U) \geq \lceil \log_4 \mathrm{Sch}(U) \rceil. \tag{6.47}$$

We will prove this result shortly. It provides a general technique for proving lower bounds on the coherent communication complexity of a given unitary operator. In order to make use of the bound, we must have a means of determining the Schmidt number of a given unitary operator. Fortunately, such a means is provided by the Schmidt decomposition, described in Appendix A.

Recall that the space of operators on the joint system AB can be regard as a Hilbert space formed from the tensor product of the Hilbert space of operators on system A with the Hilbert space of operators on system B. Any convenient inner product may be used to turn the vector spaces of operators on systems A and B into Hilbert spaces; we will use the trace inner product, $(A_1, A_2) \equiv \mathrm{tr}(A_1^\dagger A_2)$ and $(B_1, B_2) \equiv \mathrm{tr}(B_1^\dagger B_2)$. This inner product, in turn, gives rise to a Schmidt decomposition for vectors (that is, operators) on the joint space AB,

$$U = \sum_i A_i \otimes B_i, \tag{6.48}$$

where the sets A_i and B_i are guaranteed to be orthogonal with respect to the trace inner product. Moreover, properties of the Schmidt decomposition guarantee that this decomposition contains the minimal number of operators possible.

In practice, finding the Schmidt decomposition of an operator may be done in a straightforward manner, along the lines outlined in Appendix A, and we will not recap that method here in the slightly different operator language. It is, however, instructive to look at a couple of examples of the use of the Schmidt decomposition. First, we give the operator-Schmidt decomposition for the controlled not operation,

$$C = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes X. \tag{6.49}$$

Notice that this neatly encapsulates the verbal description often used for this process – if the control qubit is zero, the data qubit is left alone, while if the control qubit is one, the data qubit is flipped, while still emphasizing that this process is undertaken *coherently*, a point that may not always be clear from verbal descriptions of the controlled not gate.

To prove the lower bound (6.47), consider a general protocol for computing U using qubit communication between Alice and Bob. Let U_s be the total unitary operation performed after s steps in the protocol. Notice that $U_0 = I$ has $\mathrm{Sch}(U_0) = 1$. Once again, using the techniques of Chapter 3, we may introduce work bits to ensure that, without loss of generality, we may restrict ourselves to consideration of unitary operations. Note that the operations which may be performed during the protocol are of four types:

1. Alice does a unitary operation on her qubits.
2. Bob does a unitary operation on his qubits.
3. Alice sends a qubit to Bob.
4. Bob sends a qubit to Alice.

Clearly, $\mathrm{Sch}(U_s) = \mathrm{Sch}(U_{s+1})$ if steps 1 or 2 are carried out. Suppose $U_s = \sum_i A_{s,i} \otimes B_{s,i}$ is a minimal decomposition of U_s. Suppose Alice sends a qubit Q to Bob, leaving Alice with a system A'. Note that

$$A_{s,i} = \sum_{j=1}^{4} A'_{s,i,j} \otimes Q_{s,i,j}, \tag{6.50}$$

for some set of four operators $A'_{s,i,j}$ on A' and $Q_{s,i,j}$ on Q. Thus,

$$U_{s+1} = \sum_{i,j=1\ldots4} A'_{s,i,j} \otimes Q_{s,i,j} \otimes B_{s,i}, \tag{6.51}$$

from which we deduce that

$$\mathrm{Sch}(U_{s+1}) \leq 4\,\mathrm{Sch}(U_s). \tag{6.52}$$

Similarly, if Bob sends a qubit to Alice then $\mathrm{Sch}(U_{s+1}) \leq 4\,\mathrm{Sch}(U_s)$. Putting these observations together, if Alice and Bob employ q qubits of communication to compute U, then we must have $\mathrm{Sch}(U) \leq 4^q$, from which we have the general lower bound, (6.47).

A simple application of this lower bound is to the communication complexity of the swap operation. Suppose Alice and Bob each have n qubits, which they wish to swap. The unitary operator implementing this swap has Schmidt decomposition

$$U = \sum_{ij} |i\rangle\langle j| \otimes |j\rangle\langle i|, \tag{6.53}$$

where the sum is over all computational basis states $|i\rangle$ and $|j\rangle$. Thus $\mathrm{Sch}(U) = 4^n$, from which it follows that the coherent communication complexity is at least n qubits. The obvious method to achieve such a swap is for Alice to send her n qubits to Bob, and for Bob to send his n qubits to Alice, for a total cost of $2n$ qubits. Once again, as for the quantum Fourier transform, we are within a factor of two of knowing the exact quantum communication complexity. In actual fact, it is straightforward to adapt the methods used to prove the capacity theorem of section 6.2 to show that at least $2n$ qubits of quantum communication must be employed to perform the swap operation; this is easy, and the proof will be omitted. Nevertheless, this example of the swap operation provides a simple example where the general lower bound (6.47) provides useful information.

A second, less trivial application of (6.47) is to the quantum Fourier transform. We could explicitly work out the Schmidt decomposition for the quantum Fourier transform by using the procedure in Appendix A. However, in this case we can fortuitously note that the quantum Fourier transform on $2n$ qubits can be written

$$\begin{aligned} U &= \sum_{x_1,\ldots,x_n,y_1,\ldots,y_n} \left[\left(|0\rangle + e^{2\pi i 0.y_n}|1\rangle\right)\ldots\left(|0\rangle + e^{2\pi i 0.y_1\ldots y_n}|1\rangle\right)\langle x_1,\ldots,x_n|\right] \otimes \\ &\quad \left[\left(|0\rangle + e^{2\pi i 0.x_n y_1\ldots y_n}|1\rangle\right)\ldots\left(|0\rangle + e^{2\pi i 0.x_1\ldots x_n y_1\ldots y_n}|1\rangle\right)\langle y_1,\ldots,y_n|\right]. \end{aligned} \tag{6.54}$$

Simple algebra verifies that, as written, this is already the operator Schmidt decomposition for the quantum Fourier transform. It follows that $\mathrm{Sch}(U) = 2^{2n} = 4^n$, and thus, by (6.47), the coherent quantum communication complexity of the quantum Fourier transform is at least n qubits.

Admittedly, this is a result which we were able to prove earlier, by different means, however it is interesting to see that this result can be obtained as a special case of a more general result. To what other problems might it be possible to apply this general technique? Unfortunately, there do not seem to be many interesting unitary operations known, for quantum computation. One problem of some interest would be to investigate the quantum communication complexity of the Fourier transform over an arbitrary Abelian group rather than the group of integers modulo 2^n, as we have been considering. Nevertheless, this is a somewhat artificial problem. Less artificial is the iteration used by Grover [77] in his search algorithm; unfortunately, it is easy to see that this can be done using two qubits of communication, making the problem rather trivial.

Another problem, recently suggested to me by Raymond Laflamme, is that of evaluating the difficulty of performing quantum error correction in a distributed fashion. This is a problem which is potentially of great interest in schemes for distributed quantum computation, such as that suggested by Cirac, Zoller, Kimble and Mabuchi [43]. Laflamme has also asked me whether the above proof techniques can be adapted to a different model of distributed computation in which the allowed operation is not qubit communication between the parties, but rather quantum gates which may be performed jointly by the parties. The answer is that yes, these techniques may be adapted in a straightforward manner; a detailed working out of these developments will appear elsewhere.

6.5 A unified model for communication complexity

We have considered two broadly different classes of models for communication complexity – a class involving the computation of classical functions, using quantum resources, and a class involving the computation of quantum functions, using quantum resources. In this section a formalism is briefly outlined which has both these classes of models as special limiting cases.

An obvious means of generalizing coherent communication complexity is to consider the communication complexity for an arbitrary *quantum operation*. For example, suppose $\mathcal{E}$ is a complete quantum operation acting, jointly, on two systems, A and B. What is the minimal number of qubits

which must be communicated between A and B if the quantum operation $\mathcal{E}$ is to be implemented exactly? Another very interesting case is the performance of collective measurements on the system; suppose we have a set $\mathcal{E}_0, \mathcal{E}_1$ of incomplete quantum operations which represent a measurement. How much quantum communication must be performed between Alice and Bob if they are to be able to perform that measurement?

For definiteness, we will study the case when $\mathcal{E}$ is a complete quantum operation on AB, which is to be implemented exactly, there is no preshared resource existing between Alice and Bob, and communication is to be carried out using qubits alone. Each of the choices implied in the previous sentence could be varied to provide problems of considerable interest, but we will restrict ourselves to a single problem. Furthermore, we will not consider the very interesting problem of *families* of communication problems, which imply additional *uniformity requirements* for communication protocols, along the lines sketched in item 6 on page 27 for the quantum circuit model of quantum computation. We denote the communication complexity in this model by $Q_0(\mathcal{E})$, the minimal number of qubits that must be communicated in order to compute the quantum operation $\mathcal{E}$ *exactly*.

Let $\mathcal{E}_1$ and $\mathcal{E}_2$ be two complete quantum operations on AB. We say $\mathcal{E}_1 \leq \mathcal{E}_2$, read $\mathcal{E}_1$ *can be reduced to* $\mathcal{E}_2$, if there exist complete quantum operations $\mathcal{E}_A, \mathcal{E}'_A$ on system A and $\mathcal{E}_B, \mathcal{E}'_B$ on system B such that

$$\mathcal{E}_1 = (\mathcal{E}'_A \otimes \mathcal{E}'_B) \circ \mathcal{E}_2 \circ (\mathcal{E}_A \otimes \mathcal{E}_B). \tag{6.55}$$

It is easily verified that $\leq$ is a partial order on the set of quantum operations. It is clear that Q_0 preserves this order, $Q_0(\mathcal{E}_1) \leq Q_0(\mathcal{E}_2)$, since to perform $\mathcal{E}_1$, all we need do is perform $\mathcal{E}_A$ on system A, $\mathcal{E}_B$ on system B, then $\mathcal{E}_2$, and finish by applying $\mathcal{E}'_A$ on system A followed by $\mathcal{E}'_B$ on system B, for a total cost the same as the communication cost to compute $\mathcal{E}_2$.

This result, incidentally, is a special case of a more general *triangle inequality for communication complexity* . This is the obvious statement that if $\mathcal{E}_1$ and $\mathcal{E}_2$ are complete quantum operations then

$$Q_0(\mathcal{E}_2 \circ \mathcal{E}_1) \leq Q_0(\mathcal{E}_2) + Q_0(\mathcal{E}_1). \tag{6.56}$$

Let $f : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$ be any classical function. Define a complete quantum operation $\mathcal{E}_f$ which has as input n qubits from system A, n qubits from system B, and as output, 1 qubit system A, by the condition that $\mathcal{E}_f \equiv \mathcal{E}_1 \circ \mathcal{D}$, where $\mathcal{D}$ completes decoheres the system in the computational basis, and

$$\mathcal{E}_1(|x\rangle\langle x| \otimes |y\rangle\langle y|) \equiv |f(x,y)\rangle\langle f(x,y)|. \tag{6.57}$$

That is, $\mathcal{E}_f$ is a quantum operation which takes the state $|x\rangle|y\rangle$ as input, and outputs $|f(x,y)\rangle$, while destroying all coherences between computational basis states.

Let $\mathcal{E}$ be any quantum operation on AB which we would naturally think of as computing f. That is, we require that, on input of the state $|x,y\rangle$, the quantum operation $\mathcal{E}$ should output $|f(x,y)\rangle$ on a fixed one of Alice's qubits. Let $\mathcal{D}$ be the operation which decoheres all Alice's qubits and all Bob's qubits. The operation of decoherence in the computational basis can be performed locally by both Alice and Bob, so $\mathcal{E} \circ \mathcal{D} \leq \mathcal{E}$. Furthermore, recall from Chapter 3 that the partial trace is a complete quantum operation. This can certainly be done locally: we are merely ignoring a system! Thus, $\mathcal{E}_f \leq \mathcal{E} \circ \mathcal{D} \leq \mathcal{E}$, from which it follows that $Q_0(\mathcal{E}_f)) \leq Q_0(\mathcal{E})$. It follows that if we wish to calculate the communication complexity of the classical function f, it suffices to calculate the communication complexity of the quantum operation $\mathcal{E}_f$.

Aesthetically, this is a pleasing result; it allows us to connect the communication complexity of the classical function f to the communication complexity of a single quantum operation. Thus,

the general question of the communication complexity of a quantum operation contains both the coherent communication complexity, and the quantum communication complexity of a classical function as special cases.

6.6 Conclusion

Distributed classical computation is still only incompletely understood; how much more true this is in the quantum case! What are some of the interesting open problems and directions for research in the study of distributed quantum computation? In this section I enumerate a few of the problems which I believe are particularly interesting and important:

1. Find a general simulation technique for the entanglement assisted classical communication model introduced by Cleve and Buhrman [44] in the qubit communication model introduced by Yao [198].

2. How is the coherent communication complexity affected by the presence of one or more of the following resources: pre-shared entanglement, pre-shared classical correlation, or classical communication?

3. What are the coherent communication complexities for some more, truly quantum, operations, beyond the quantum Fourier transform?

4. How are results on quantum communication complexity affected by the presence of noise in the communications channel?

5. Can notions of quantum communication complexity be used to define measures of entanglement in multipartite quantum systems? In Chapter 8 we will study quantitative measures of the the entanglement between two quantum systems. These measures are based upon *resource problems*. Quantum communication complexity is a natural source of such resource problems, and it is possible that one of these resource problems may be used to provide a good measure of entanglement, perhaps even for systems consisting of more than two parts, a major bugbear of present efforts to study entanglement.

This is just a sample of the sorts of questions which naturally arise out of consideration of distributed quantum computation. Judging from the rapid progress over the past eighteen months, I expect that the field of quantum communication complexity will be one of the major areas of significant development in quantum information theory over the next few years. This progress, in turn, should help stimulate other parts of the field with new insights into the nature of quantum information.

Summary of Chapter 6: Quantum communication complexity

- **Capacity theorem for communication using qubits:** Suppose that Alice possesses n bits of information, and wants to convey this information to Bob. Suppose that Alice and Bob possess no prior entanglement but qubit communication in either direction is allowed. Let n_{AB} be the number of qubits Alice sends to Bob, and n_{BA} the number of qubits Bob sends to Alice. Then, Bob can acquire the n bits if and only if the following inequalities are satisfied:

$$\begin{aligned} n_{AB}, n_{BA} &\geq 0 && (6.58) \\ n_{AB} &\geq \lceil n/2 \rceil && (6.59) \\ n_{AB} + n_{BA} &\geq n. && (6.60) \end{aligned}$$

- **Entanglement-assisted communication complexity:** The number of bits of classical information that must be communicated between Alice and Bob if they are to compute a given (classical) function, given that they may preshare an arbitrary entanglement. In order to compute the inner product, modulo 2, of an n bit string belonging to Alice, and an n bit string belonging to Bob, requires precisely n bits of classical communication.

- **Coherent quantum communication complexity:** How many qubits need to be communicated between Alice and Bob if they are to compute a quantum function, that is, some family of quantum operations? For the $2n$ qubit quantum Fourier transform, n qubits belonging to Alice and n qubits belonging to Bob, an n qubit lower bound can be proved. Furthermore, to do an approximate Fourier transform a distance $D < 1/2$ from the exact Fourier transform, at least $n(1-D) - h(D)$ qubits must be sent, where $h(\cdot)$ is the binary entropy.

- **Lower bound on the coherent communication complexity:**

$$Q_0(U) \geq \lceil \log_4 \mathrm{Sch}(U) \rceil. \quad (6.61)$$

Chapter 7

Quantum data compression

The storage of states produced by a quantum source using the fewest possible resources is a fundamental problem of quantum information theory. Schumacher and co-workers [154, 94, 8] have shown that a quantum source ρ produced by picking from an ensemble of quantum states $\{p_i, |\psi_i\rangle\}$ may be compressed so that it requires only $S(\rho)$ qubits per source state for reliable storage. Barnum, Fuchs, Jozsa and Schumacher [8] have shown that $S(\rho)$ qubits per source state is the minimal resources required for reliable storage. The basic idea of quantum data compression is illustrated in figure 7.1. A quantum source ρ on d qubits is used n times. A compression operation, $\mathcal{C}$ is used to compress that source into roughly $nS(\rho)$ qubits. At some later time, a decompression operation, $\mathcal{D}$ is used to recover the original state produced by the source, with high fidelity.

This Chapter addresses the compression of a quantum source producing states ρ which are entangled with another, inaccessible quantum system, using the tools introduced in Chapter 5, especially the *dynamic fidelity*. As motivation for the use of the dynamic fidelity, we might imagine that we are trying to compress part of the memory of a quantum computer, and that we wish to recover the entanglement with the rest of the quantum computer at some later time. This approach based upon the dynamic fidelity is quite different to the work of Schumacher and collaborators, who used a different measure of reliability, to be discussed below. An advantage of the present approach is that several of the proofs appear more natural than in the approach pioneered by Schumacher [154]. In particular, Schumacher did not find a simple proof that $S(\rho)$ is the minimal resources required to do quantum data compression. Schumacher [154] and later Schumacher and Jozsa [94] gave incomplete proofs of this result that did not consider the most general possible decoding schemes.

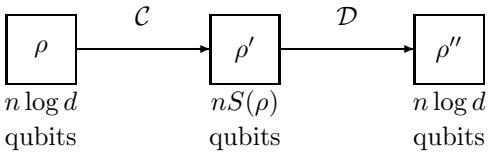

Figure 7.1: Quantum data compression. The compression operation $\mathcal{C}$ compresses a quantum source ρ stored in $n \log d$ qubits into $nS(\rho)$ qubits. The source is accurately recovered via the decompression operation $\mathcal{D}$.

The proof was later completed by Barnum *et al* [8], using an ingenious but rather complicated argument. The techniques introduced in this Chapter lead to simple and direct proofs.

Using these techniques we also study the problem of *universal quantum data compression*. In classical information theory the existence of universal compression algorithms, such as the well-known Lempel-Ziv [201] algorithm, is an important and useful fact, exploited in many widely available programs and devices, such as the UNIX *compress* program. A universal compression algorithm is one which can compress a large class of sources, not just a single source. We study the limits to universal quantum data compression, and exhibit a quantum scheme which is universal with respect to a large set of sources.

The Chapter is organized as follows. Section 7.1 defines the *Schmidt number* of a pure state of a composite system, and studies some of the properties of the Schmidt number. Many of the results in the Chapter are simple counting arguments based on properties of the Schmidt number. Section 7.2 reviews some basic facts about typical subspaces of a quantum source. With these tools in hand, section 7.3 proves the quantum data compression theorem. Section 7.4 extends this result to the case of a quantum channel with a side channel for classical information. Section 7.5 studies universal quantum data compression. Section 7.6 concludes the Chapter with a discussion of future directions. The work reported in the Chapter is largely my original work, with the exception of Section 7.2, which is based upon the work of Schumacher and Jozsa.

7.1 Schmidt numbers

This section reintroduces an extremely useful tool for proving properties about entangled systems: the *Schmidt number* of an entangled pure state. This tool was also used, in a different and less central guise, in Chapter 6, to prove results about quantum communication complexity. All of the results in this section are rather elementary, yet they play a crucial role in our proof of the quantum data compression theorem.

For convenience, we restate the *Schmidt decomposition* theorem, an extremely useful structural theorem for pure states of composite quantum systems, proved in Appendix A. Suppose $|AB\rangle$ is a pure state for some joint system AB. Then there exists a *Schmidt decomposition* for $|AB\rangle$,

$$|AB\rangle = \sum_i \sqrt{p_i}|i^A\rangle|i^B\rangle \tag{7.1}$$

where $|i^A\rangle$ is an orthonormal basis for A and $|i^B\rangle$ is an orthonormal basis for B and $p_i \geq 0, \sum_i p_i = 1$. Recall also that these bases are identical to the bases in which the reduced density operators on A and B are diagonal, since

$$A \equiv \mathrm{tr}_B(|AB\rangle\langle AB|) = \sum_i p_i |i^A\rangle\langle i^A| \tag{7.2}$$

$$B \equiv \mathrm{tr}_A(|AB\rangle\langle AB|) = \sum_i p_i |i^B\rangle\langle i^B|. \tag{7.3}$$

We define the *Schmidt number* of $|AB\rangle$ to be the number of non-zero p_i in the Schmidt decomposition of $|AB\rangle$. An equivalent and rather useful way of defining the Schmidt number is using the concept of *support*. Given a diagonalizable operator D the *support* $\mathrm{supp}(D)$ of that operator is defined to be the vector space spanned by those eigenvectors of D whose corresponding eigenvalue is non-zero. Clearly, the Schmidt number is equal to the dimension of the support of A, which is also equal to the dimension of the support of B,

$$\mathrm{Sch}(|AB\rangle) \equiv \dim \mathrm{supp}(A) = \dim \mathrm{supp}(B). \tag{7.4}$$

The Schmidt number has two especially useful properties under operations:

1. Suppose $\mathcal{C} : Q \to Q'$ is any quantum operation mapping density operators on the space Q to density operators on the space Q'. Let RQ be a given density operator on the composite system RQ. Let $\mathcal{I}$ be the identity operation on R. Then the state of the system after the action of $\mathcal{C}$,

$$R'Q' \equiv \frac{(\mathcal{I} \otimes \mathcal{C})(RQ)}{\mathrm{tr}\left[(\mathcal{I} \otimes \mathcal{C})(RQ)\right]}, \tag{7.5}$$

can be written in the form

$$R'Q' = \sum_j q_j |R'Q'_j\rangle\langle R'Q'_j|, \tag{7.6}$$

where $q_j \geq 0, \sum_j q_j = 1$, the $|R'Q'_j\rangle$ form an orthonormal set, and

$$\mathrm{Sch}(|R'Q'_j\rangle) \leq \dim(Q'). \tag{7.7}$$

2. Let $\mathcal{D} : Q' \to Q''$ be any quantum operation mapping density operators on the space Q' to density operators on the space Q''. Let $R'Q'$ be any state of $R'Q'$. Then

$$R''Q'' \equiv \frac{(\mathcal{I} \otimes \mathcal{D})(R'Q')}{\mathrm{tr}(\mathcal{I} \otimes \mathcal{D})(R'Q')} \tag{7.8}$$

can be written in the form

$$R''Q'' = \sum_k s_k |R''Q''_k\rangle\langle R''Q''_k|, \tag{7.9}$$

where $s_k \geq 0, \sum_k s_k = 1$, the $|R''Q''_k\rangle$ are pure states, not necessarily orthonormal, and

$$\mathrm{Sch}(|R''Q''_k\rangle) \leq \dim(Q'). \tag{7.10}$$

Both properties follow immediately from the definition of the Schmidt number. Heuristically, the first property is just the obvious fact that the output ensemble from a quantum operation can not have a Schmidt number higher than the dimension of the output space. The second result is only slightly less obvious, stating that a quantum operation can not increase the Schmidt number of elements in an ensemble. Actually, it is clear that stronger results than 2 are true, and rather interesting, although we will not need such a result here. For completeness, we describe one such result:

Let $\mathcal{D} : Q' \to Q''$ be any quantum operation, and $R'Q' = |R'Q'\rangle\langle R'Q'|$ be any pure state of $R'Q'$. Defining $R''Q''$ as before, it follows that $R''Q''$ can be written in the form

$$R''Q'' = \sum_k s_k |R''Q''_k\rangle\langle R''Q''_k|, \tag{7.11}$$

where $s_k \geq 0$, $\sum_k s_k = 1$ and

$$\mathrm{Sch}(|R''Q''_k\rangle) \leq \mathrm{Sch}(|R'Q'\rangle). \tag{7.12}$$

Property 2, above, is clearly a consequence of this result, which is also immediate from the definition of the Schmidt number.

7.2 Typical subspaces

The notion of a *typical subspace* was introduced by Schumacher [154] and Jozsa and Schumacher [94] as the quantum analogue of an important notion in classical information theory, that of a typical sequence. They proved the following two results, which we will refer to jointly as the *typical subspace theorem*. Both parts of the typical subspace theorem follow easily from the weak law of large numbers, proved in the box on page 135.

Theorem 20 (Typical subspace theorem) *[154, 94]*

1. *Fix a quantum state ρ in a state space of d qubits[1]. Let $\epsilon > 0$ be given. For all n sufficiently large there exists a projector P_ϵ^n onto a space of at most $2^{n(S(\rho)+\epsilon)}$ dimensions such that*

$$tr(\rho^{\otimes n} P_\epsilon^n) > 1 - \epsilon \tag{7.13}$$

$$\left[\rho^{\otimes n}, P_\epsilon^n\right] = 0. \tag{7.14}$$

2. *Fix ρ. Let $\epsilon > 0$ be given. Let P_n be any sequence of projectors such that P_n projects onto a space of at most 2^{nR} dimensions. Then*

$$tr(\rho^{\otimes n} P_n) \leq 2^{-n(S(\rho)-R-\epsilon)} + \epsilon \tag{7.15}$$

for all sufficiently large values of n.

Proof [154, 94]

For completeness, we outline the construction of the projector P_ϵ^n onto the typical subspace. Suppose ρ has orthogonal decomposition,

$$\rho^{\otimes n} = \sum_x p_x |x\rangle\langle x|. \tag{7.16}$$

Then $\rho^{\otimes n}$ has orthogonal decomposition

$$\rho = \sum_{\mathbf{x}} p_{\mathbf{x}} |\mathbf{x}\rangle\langle\mathbf{x}|, \tag{7.17}$$

where the sum is over all sequences $\mathbf{x} = x_1, \ldots, x_n$, $p_{\mathbf{x}} \equiv p_{x_1} p_{x_2} \ldots p_{x_n}$ and $|\mathbf{x}\rangle = |x_1\rangle|x_2\rangle \ldots |x_n\rangle$. We say a sequence $\mathbf{x}$ is ϵ-typical if

$$2^{-n(S(\rho)+\epsilon)} \leq p_{\mathbf{x}} \leq 2^{-n(S(\rho)-\epsilon)}. \tag{7.18}$$

Intuitively, the sequence $\mathbf{x}$ can be thought of as the sequence of outputs produced by a classical source producing independent random variables, identically distributed according to p_x. Using the law of large numbers and taking the logarithm of the above definition, we see that in the limit as n goes to infinity, a typical sequence occurs with probability going to one. Furthermore, since the sum of a set of probabilities is at most one, and

$$p_{\mathbf{x}} \geq 2^{-n(S(\rho)+\epsilon)}, \tag{7.19}$$

we see that there at most $2^{n(S(\rho)+\epsilon)}$ ϵ-typical sequences. We now define

$$P_\epsilon^n \equiv \sum_{\epsilon-\text{typical } \mathbf{x}} |\mathbf{x}\rangle\langle\mathbf{x}| \tag{7.20}$$

[1] *The restriction to qubit systems is not necessary, but it may help make the discussion more concrete.*

to be the projector onto the ϵ-typical subspace. Note that

$$\mathrm{tr}(\rho^{\otimes n} P_\epsilon^n) = \sum_{\epsilon-\text{typical } \mathbf{x}} p_{\mathbf{x}}, \tag{7.21}$$

so by the law of large numbers, for sufficiently large n, $\mathrm{tr}(\rho^{\otimes n} P_\epsilon^n) > 1 - \delta$, for any $\delta > 0$. Setting $\delta \equiv \epsilon$ proves equation (7.13). Furthermore, by definition P_ϵ^n is diagonal in the same basis as $\rho^{\otimes n}$, and thus commutes with $\rho^{\otimes n}$, equation (7.14).

The second property now follows from the identity

$$\mathrm{tr}(\rho^{\otimes n} P_n) = \mathrm{tr}(\rho^{\otimes n} P_\epsilon^n P_n) + \mathrm{tr}(\rho^{\otimes n}(I - P_\epsilon^n) P_n) \tag{7.22}$$

and the observations that

$$\mathrm{tr}(\rho^{\otimes n} P_\epsilon^n P_n) \leq 2^{-n(S(\rho)-\epsilon)} 2^{nR}, \tag{7.23}$$

and

$$\begin{aligned} \mathrm{tr}(\rho^{\otimes n}(I - P_\epsilon^n) P_n) &\leq \mathrm{tr}(\rho^{\otimes n}(I - P_\epsilon^n)) && (7.24) \\ &\leq \epsilon. && (7.25) \end{aligned}$$

QED

The law of large numbers

Suppose we repeat an experiment a large number of times, each time measuring the value of some parameter, X. We label the results of the experiments $X_1, X_2, \ldots$. Assuming that the results of the experiments are independent, we intuitively expect that the value of the estimator

$$S_n \equiv \sum_{i=1}^{n} \frac{X_i}{n} \tag{7.26}$$

of the average $\mathbf{E}(X)$, should approach $\mathbf{E}(X)$ as $n \to \infty$. The *law of large numbers* [76] is a rigorous statement of this intuition.

Theorem (Law of large numbers) Suppose $X_1, X_2, \ldots$ are independent and identically distributed random variables, with finite first and second moments, $|\mathbf{E}(X_1)| < \infty$ and $\mathbf{E}(X_1^2) < \infty$. Then for any $\epsilon > 0$, $p(|S_n - \mathbf{E}(X)| > \epsilon) \to 0$ as $n \to \infty$.

Proof:

To begin we assume that $\mathbf{E}(X_i) = 0$. We will discuss what happens when $\mathbf{E}(X_i) \neq 0$ upon completion of the proof. Since the random variables are independent with mean zero, it follows that $\mathbf{E}(X_i X_j) = \mathbf{E}(X_i)\mathbf{E}(X_j) = 0$ when $i \neq j$, and thus

$$\mathbf{E}(S_n^2) = \frac{\sum_{i,j=1}^{n} \mathbf{E}(X_i X_j)}{n^2} = \frac{\sum_{i=1}^{n} \mathbf{E}(X_i^2)}{n^2} = \frac{\mathbf{E}(X_1^2)}{n}, \tag{7.27}$$

where the final equality follows from the fact that $X_1, \ldots, X_n$ are identically distributed. By the same token, from the definition of the expectation we have

$$\mathbf{E}(S_n^2) \quad = \quad \int dP\, S_n^2, \tag{7.28}$$

where dP is the underlying probability measure. It is clear that either $|S_n| \leq \epsilon$ or $|S_n| > \epsilon$, so we can split this integral into two pieces, and then drop one of these pieces, observing that it is non-negative,

$$\mathbf{E}(S_n^2) \quad = \quad \int_{|S_n| \leq \epsilon} dP\, S_n^2 + \int_{|S_n| > \epsilon} dP\, S_n^2 \;\geq\; \int_{|S_n| > \epsilon} dP\, S_n^2. \tag{7.29}$$

In the region of integration $S_n^2 > \epsilon^2$, and thus

$$\mathbf{E}(S_n^2) \quad \geq \quad \epsilon^2 \int_{|S_n| > \epsilon} dP \;=\; \epsilon^2 p(|S_n| > \epsilon). \tag{7.30}$$

Comparing this inequality with (7.27) we see that

$$p(|S_n| > \epsilon) \leq \frac{\mathbf{E}(X_1^2)}{n\epsilon^2}. \tag{7.31}$$

Letting $n \to \infty$ completes the proof. In the case when $\mathbf{E}(X_1) \neq 0$, it is easy to obtain the result, by defining $Y_i \equiv X_i - \mathbf{E}(X_1)$. The Y_i are a sequence of independent, identically distributed random variables with $\mathbf{E}(Y_1) = 0$ and $\mathbf{E}(Y_1^2) < \infty$. The result follows from the earlier reasoning.

QED

7.3 Quantum data compression theorem

We now have all the tools necessary to prove the quantum data compression theorem. To understand the result, we first need to make more formal the notions of quantum sources, and block encoding and decoding.

An *i.i.d. (independent, identically distributed) quantum source* $\{H, \rho\}$ consists of a Hilbert space H and a density operator ρ on that Hilbert space. The *n-blocked* source is the pair $\{H^{\otimes n}, \rho^{\otimes n}\}$. A *compression scheme of rate* R for a source $\{H, \rho\}$ is a sequence $\{\mathcal{C}^n, \mathcal{D}^n\}$ of quantum operations such that the *encoding operation* $\mathcal{C}^n$ maps the n-blocked source space $H^{\otimes n}$ into a Hilbert space H_c^n of dimension 2^{nR}, and the *decoding operation* $\mathcal{D}^n$ maps the 2^{nR} dimensional Hilbert space H_c^n back into the source Hilbert space $H^{\otimes n}$.

Our criterion for whether or not the compression-decompression procedure has been successfully accomplished is whether or not the dynamic fidelity for the total procedure is close to one. As we saw in Chapter 5, a dynamic fidelity close to one is equivalent to the requirement that the source and any entanglement it has with other systems has been well preserved by the process.

More precisely, we say that a compression scheme is *reliable* if

$$\lim_{n\to\infty} F(\rho^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) = 1. \tag{7.32}$$

A compression scheme is said to be *weakly unreliable* if it is not reliable. A compression scheme is said to be *strongly unreliable* if

$$\lim_{n\to\infty} F(\rho^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) = 0. \tag{7.33}$$

Clearly a strongly unreliable compression scheme is also weakly unreliable.

Theorem 21 (Quantum entanglement compression theorem)
Let $\{H, \rho\}$ be a quantum source.

1. *(Achievability)*
   If $R > S(\rho)$ then there exists a reliable compression scheme of rate R for the source $\{H, \rho\}$.
2. *(Weak converse).*
   If $R < S(\rho)$ then all compression schemes are weakly unreliable.
3. *(Strong converse)*
   If $R < S(\rho)$ then all compression schemes are strongly unreliable.

Obviously the strong converse implies the weak converse. Both results are stated here because we will give a proof of the weak converse which is independent of the proof of the strong converse.

Proof

Proof of achievability

The compression scheme used to prove achievability is exactly the same as that used by Jozsa and Schumacher [94], although the analysis is made slightly different by the use of dynamic fidelity as the reliability criterion.

Let $\epsilon > 0$ be such that $S(\rho) + \epsilon \leq R$. Define P_ϵ^n to be the projector onto the ϵ-typical subspace, and use T_ϵ^n to denote the ϵ-typical subspace. By the typical subspace theorem, for all n sufficiently large,

$$\mathrm{tr}(\rho^{\otimes n} P_\epsilon^n) \geq 1 - \epsilon, \tag{7.34}$$

and

$$\dim(T_\epsilon^n) \leq 2^{nR}. \tag{7.35}$$

Let H_c^n be any 2^{nR} dimensional Hilbert space containing T_ϵ^n. The encoding is done in the following fashion. First a measurement is made, described by the complete set of orthogonal projectors $P_\epsilon^n, I - P_\epsilon^n$, with corresponding outcomes we will call 1 and 0. If outcome 0 occurs nothing more is done and the state is left in the typical subspace. If outcome 1 occurs then we replace the state of the system with some standard state $|0\rangle$ chosen from the typical subspace. It follows that the encoding is a map

$$\mathcal{C}^n : H^{\otimes n} \rightarrow H_c^n \tag{7.36}$$

and has the operator-sum representation

$$\mathcal{C}^n(\sigma) \equiv P_\epsilon^n \rho P_\epsilon^n + \sum_i A_i \sigma A_i^\dagger, \tag{7.37}$$

where

$$A_i = |0\rangle\langle i| \tag{7.38}$$

and $|i\rangle$ is an orthonormal basis for the orthocomplement of the typical subspace.

The decoding operation

$$\mathcal{D}^n : H_c^n \rightarrow H^{\otimes n} \tag{7.39}$$

is just the identity on H_c^n, $\mathcal{D}^n(\sigma) = \sigma$. With these definitions for the encoding and decoding it follows that

$$\begin{aligned} F(\rho^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) &= |\mathrm{tr}(\rho^{\otimes n} P_\epsilon^n)|^2 + \sum_i |\mathrm{tr}(\rho^{\otimes n} A_i)|^2 & (7.40) \\ &\geq |\mathrm{tr}\rho^{\otimes n} P_\epsilon^n)|^2 & (7.41) \\ &\geq |1-\epsilon|^2 \geq 1 - 2\epsilon, & (7.42) \end{aligned}$$

where the last line follows from the theorem of typical subspaces. But ϵ can be made arbitrarily small and thus it follows that there exists a reliable compression scheme $\{\mathcal{C}^n, \mathcal{D}^n\}$ of rate R whenever $S(\rho) < R$.

Proof of the weak converse

In Chapter 10, on page 194, we prove the following result, known as the *entropy-fidelity* inequality. For any ρ and complete quantum operations $\mathcal{C}$ and $\mathcal{D}$,

$$S(\rho) \leq I(\rho, \mathcal{C}) + 2 + 4(1 - F(\rho, \mathcal{D} \circ \mathcal{C})) \log d, \tag{7.43}$$

where d is the dimension of the Hilbert space of ρ and $I(\rho, \mathcal{C})$ is the *coherent information* defined by

$$I(\rho, \mathcal{C}) \equiv S(\mathcal{C}(\rho)) - S(\rho, \mathcal{C}), \tag{7.44}$$

and $S(\rho, \mathcal{C})$ is a non-negative quantity known as the *entropy exchange*. From the non-negativity of the entropy exchange it follows that $I(\rho, \mathcal{C}) \leq S(\mathcal{C}(\rho))$ and thus from the entropy-fidelity inequality

$$S(\rho) \leq S(\mathcal{C}(\rho)) + 2 + 4(1 - F(\rho, \mathcal{D} \circ \mathcal{C})) \log d. \tag{7.45}$$

Applying (7.45) to $\rho^{\otimes n}, \mathcal{C}^n, \mathcal{D}^n$ and noting that $S(\mathcal{C}^n(\rho^{\otimes n})) \leq \log 2^{nR} = nR$ we see that

$$nS(\rho) \leq nR + 2 + 4n(1 - F(\rho, \mathcal{D}^n \circ \mathcal{C}^n)) \log d, \tag{7.46}$$

where d is the dimension of the source space H. Dividing by n and taking the limit as $n \to \infty$ we see that

$$S(\rho) \leq R + \lim_{n\to\infty} (1 - F(\rho, \mathcal{D}^n \circ \mathcal{C}^n)) \log d \tag{7.47}$$

(when the limit exists). Thus, for reliable transmission we obtain

$$S(\rho) \leq R. \tag{7.48}$$

It follows that if $R < S(\rho)$ then all compression schemes must be weakly unreliable.

Proof of the strong converse

Suppose $|RQ\rangle$ purifies ρ. Then taking n copies of RQ, $|RQ_n\rangle \equiv |RQ\rangle^{\otimes n}$ purifies $\rho^{\otimes n}$. Define

$$\rho_n^{RQ} \equiv |RQ_n\rangle\langle RQ_n| \tag{7.49}$$
$$\rho_n^{RQ'} \equiv (\mathcal{I} \otimes \mathcal{C}^n)(\rho_n^{RQ}) \tag{7.50}$$
$$\rho_n^{RQ''} \equiv (\mathcal{I} \otimes \mathcal{D}^n)(\rho_n^{RQ'}), \tag{7.51}$$

where $\mathcal{I}$ is the identity operation on $R^{\otimes n}$.

From properties 1 and 2 of the Schmidt number enumerated on page 132 we see that $\rho_n^{RQ''}$ can be written in the form

$$\rho_n^{RQ''} = \sum_i p_i |\phi_i^{RQ''}\rangle\langle\phi_i^{RQ''}|, \tag{7.52}$$

where $p_i \geq 0, \sum_i p_i = 1$ and

$$\mathrm{Sch}(|\phi_i^{RQ''}\rangle) \leq 2^{nR}. \tag{7.53}$$

By definition we have

$$F(\rho^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) = \langle RQ_n|\rho_n^{RQ''}|RQ_n\rangle, \tag{7.54}$$

and thus

$$F(\rho^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) = \sum_i p_i |\langle RQ_n|\phi_i^{RQ''}\rangle|^2. \tag{7.55}$$

To bound the dynamic fidelity we examine the individual terms in this equation. Note that

$$|\langle RQ_n|\phi_i^{RQ''}\rangle|^2 = \mathrm{tr}\left[\rho_n^{RQ}\sigma_i^{RQ''}\right], \tag{7.56}$$

where $\sigma_i^{RQ''} \equiv |\phi_i^{RQ''}\rangle\langle\phi_i^{RQ''}|$. Let P_i be the projector onto the support of $\sigma_i^{Q''}$. Notice that

$$\sigma_i^{Q''} \leq (I^R \otimes P_i), \tag{7.57}$$

as an operator inequality, and thus

$$\mathrm{tr}\left[\rho^{RQ}\sigma_i^{RQ''}\right] \leq \mathrm{tr}\left[\rho_n^Q P_i\right]. \tag{7.58}$$

But

$$\dim(P_i) \quad = \quad \mathrm{Sch}(|\phi_i^{RQ''}\rangle) \leq 2^{nR} \tag{7.59}$$

and by the second part of the typical subspace theorem it follows that for sufficiently large n,

$$\mathrm{tr}\left[\rho_n^Q P_i\right] \leq 2^{-n(S(\rho)-R-\epsilon)}. \tag{7.60}$$

Putting it all together we have

$$|\langle RQ_n|\phi_i^{RQ''}\rangle|^2 \leq 2^{-n(S(\rho)-R-\epsilon)} + \epsilon, \tag{7.61}$$

for sufficiently large n. Note, incidentally, that how large n needs to be depends only on ϵ and R, and not on i. Inserting this equation into (7.55) gives the result

$$F(\rho^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) \leq 2^{-n(S(\rho)-R-\epsilon)} + \epsilon \tag{7.62}$$

for all $\epsilon > 0$, for sufficiently large n. It follows that if $R < S(\rho)$ then

$$F(\rho^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) \to 0 \tag{7.63}$$

as $n \to \infty$, which is what we set out to prove.

QED

Note that the proof of the weak converse presented here is more difficult than that of the strong converse, since the proof makes implicit use of the strong subadditivity inequality for entropies via the entropy fidelity inequality, (7.43), which is proved in Chapter 10. Nevertheless, the proofs of both converse theorems appear quite natural compared to the proof found in [8]. The proof of the strong converse, especially, depends only upon elementary facts. The same proof of the weak converse was obtained independently by Allahverdyan and Saakian [3]. In the next section similar ideas will be used to prove a stronger version of the weak converse than was proved by Allahverdyan and Saakian.

7.4 Quantum data compression with a classical side channel

In general, measurements are performed during compression of a quantum source. These measurements yield classical information: the measurement outcome. Suppose a *classical side-channel* is available, as shown in figure 7.2, so that the outcomes of the measurements performed during compression are available to assist in decompression.

Intuitively, it seems likely that such a classical side channel cannot assist in the compression of the quantum information. The reason is because any measurement performed on the quantum system cannot obtain information about that state without causing some irreversible disturbance to the state. Thus, for the storage to be reliable it is necessary that the classical side channel contain no information about the source. Therefore, it seems unlikely that a classical side channel can assist in compression.

Such intuitive arguments are at best heuristic guides. Consider that quantum teleportation, review in Chapter 2, involves the use of a classical side channel which is necessary for the recovery of the quantum state of interest. Yet the information in that side channel contains no information about the state which is being teleported. This section gives a rigorous proof that the use of a classical side channel does not decrease the minimal storage requirements for a quantum state.

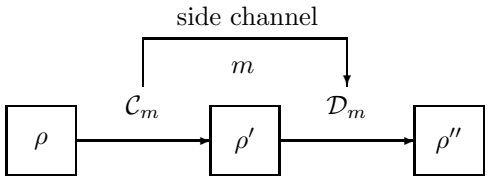

Figure 7.2: Quantum data compression with a classical side channel. Results of measurements made during compression are made available to assist during decompression.

A *compression scheme of rate R with side information* for a source $\{H, \rho\}$ is a sequence $\{\mathcal{C}_i^n, \mathcal{D}_i^n\}$ of quantum operations such that the *encoding operations* $\mathcal{C}_i^n$ map the n-blocked source space $H^{\otimes n}$ into a Hilbert space H_c^n of dimension 2^{nR}, $\sum_i \mathcal{C}_i^n$ is a complete quantum operation, and the *decoding operation* $\mathcal{D}_i^n$ is a complete quantum operation that maps the 2^{nR} dimensional Hilbert space H_c^n back into the source Hilbert space $H^{\otimes n}$.

Such a compression scheme is *reliable* if

$$\lim_{n\to\infty} F(\rho^{\otimes n}, \sum_i \mathcal{D}_i^n \circ \mathcal{C}_i^n) = 1, \tag{7.64}$$

and is said to be *weakly unreliable* if it is not reliable. Such a compression scheme is said to be *strongly unreliable* if

$$\lim_{n\to\infty} F(\rho^{\otimes n}, \sum_i \mathcal{D}_i^n \circ \mathcal{C}_i^n) = 0. \tag{7.65}$$

Clearly a strongly unreliable compression scheme is also weakly unreliable.

Theorem 22 (Quantum entanglement compression with a classical side-channel)
Let $\{H, \rho\}$ be a quantum source. Then:

1. *(Weak converse).*
   If $R < S(\rho)$ then all compression schemes with side information are weakly unreliable.
2. *(Strong converse)*
   If $R < S(\rho)$ then all compression schemes with side information are strongly unreliable.

Achievability need not be considered, since the compression scheme in the last section already achieves the best possible rate of compression without the need for classical side information. Once again the strong converse implies the weak converse, and we outline independent proofs of the two results.

Proof

Outline proof of the weak converse:

We use the "generalized entropy-fidelity" lemma from section 10.7.1. This result states that for any set of operations $\mathcal{C}_i, \mathcal{D}_i$ such that $\sum_i \mathcal{C}_i$ and $\mathcal{D}_i$ are trace-preserving,

$$\begin{aligned} S(\rho) \quad \leq \quad & \sum_i \mathrm{tr}(\mathcal{C}_i(\rho)) I(\rho, \mathcal{C}_i) + 2 + \\ & 4(1 - F(\rho, \sum_m \mathcal{D}_i \circ \mathcal{C}_i)) \log d. \end{aligned} \tag{7.66}$$

Applying this result with $\rho = \rho^{\otimes n}$, $d = d^n$, $\mathcal{C}_i = \mathcal{C}_i^n$, and $\mathcal{D}_i = \mathcal{D}_i^n$ gives the result by the same reasoning used earlier in the proof of the weak converse without a classical side channel.

Outline proof of the strong converse:

Define

$$\rho_n^{RQ''} \equiv \left[\mathcal{I} \otimes \left(\sum_i \mathcal{D}_i \circ \mathcal{C}_i\right)\right](\rho_n^{RQ}). \tag{7.67}$$

The proof of the strong converse in the presence of a classical side channel now proceeds exactly as that for the strong converse without a classical side channel.

QED

7.5 Universal data compression with a classical side channel

In order to do quantum data compression by the means we have described it is necessary to know the source density operator ρ in order to construct the projector onto the typical subspace of ρ. An elegant quantum circuit which does essentially this has been described by Cleve and DiVincenzo [45], demonstrating the possibility of doing quantum data compression on a quantum computer.

In order for such methods for data compression to work it is necessary to know the source density operator. Many popular algorithms for data compression on classical computers work in a very different way, succeeding for all source distributions within some large class of possible source distribution. Essentially, they do this by sampling from the source to build up a good knowledge of the nature of the source distribution, and use that knowledge to perform encoding in an appropriate manner. For example, Lempel-Ziv coding [201], used in popular programs such as the UNIX *compress*, compresses all stationary ergodic sources to the limit allowed by Shannon's noiseless channel coding theorem. Obviously, such algorithms are highly desirable whenever one does not know the source distribution *a priori*.

Clearly, it is desirable to have analogous universal *quantum* data compression algorithms. At first sight this may appear hopeless: classically, universal data compression works because the compressor obtains information about a source as the source is sampled, which asymptotically allows efficient compression. Quantum mechanically, we know that we cannot obtain information about a source without disturbing it, so it would seem difficult to perform universal quantum data compression. Despite, this difficulty, we will see in this section that it is possible to perform a useful form of universal quantum data compression, making use of an auxiliary classical side channel. We begin with a simple example.

Suppose ρ_1 and ρ_2 are density operators such that $S(\rho_1) < S(\rho_2)$. Suppose P_i^n is the projector onto the ϵ-typical subspace of $\rho_i^{\otimes n}$ and define

$$Q_i^n \equiv I^{\otimes n} - P_i^n, \tag{7.68}$$

the projector onto the complement of the ϵ-typical subspace of $\rho_i^{\otimes n}$. The data compression stage is done as follows:

1. Perform the measurement defined by the projectors P_1^n and Q_1^n. Label the corresponding measurement result M_1. $M_1 = 0$ if P_1^n occurred, and $M_1 = 1$ if Q_1^n occurred.

2. If $M_1 = 1$, then perform the measurement defined by the projectors P_2^n and Q_2^n. Label the corresponding measurement result M_2. $M_2 = 0$ if P_2^n occurred, and $M_2 = 1$ if Q_2^n occurred. If $M_1 = 0$ then define $M_2 = 1$.

3. If $M_1 = 1$ and $M_2 = 1$ then set $M_3 \equiv 0$.

Let r be the minimal value such that $M_r = 0$. Suppose $r = 3$. Then we don't send anything; we give up. If $r = 1$ then we send (a unitary transformation of) the typical subspace of ρ_1, which asymptotically requires $S(\rho_1)$ qubits per use of the source. If $r = 2$ then we send (a unitary transformation of) the typical subspace of ρ_2, which asymptotically requires $S(\rho_2)$ qubits per use of the source. The decoding operation is merely to map the state of the transmitted qubits back into the appropriate typical subspace for ρ_1 or ρ_2, depending on whether $r = 1$ or $r = 2$. If $r = 3$, then no decoding is performed; we give up.

If the source is ρ_1, then the dynamic fidelity for this compression scheme satisfies

$$F(\rho_1^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}) \geq |\mathrm{tr}(P_1^n \rho_1^{\otimes n})|^2, \tag{7.69}$$

which, as we saw earlier, asymptotically tends to 1 as $n \to \infty$. Much more interesting is the case when ρ_2 is the source. In this case the dynamic fidelity satisfies

$$F(\rho_2^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) \geq |\mathrm{tr}(P_2^n Q_1^n \rho_2^{\otimes n})|^2. \tag{7.70}$$

Note, however, that

$$|\mathrm{tr}(P_2^n P_1^n \rho_2^{\otimes n})| \leq \mathrm{tr}(P_1^n |\rho_2^{\otimes n} P_2^n|). \tag{7.71}$$

But the largest eigenvalue of $|\rho_2^{\otimes n} P_2^n|$ is less than the largest eigenvalue of $\rho_2^{\otimes n}$, so by the theorem of typical subspaces

$$|\mathrm{tr}(P_2^n P_1^n \rho_2^{\otimes n})| \leq 2^{n(S(\rho_1)+\epsilon)} 2^{-n(S(\rho_2)-\epsilon)} = 2^{n(S(\rho_1)-S(\rho_2)+2\epsilon)}, \tag{7.72}$$

which tends to zero as $n \to \infty$ and we allow ϵ to tend to zero. Thus

$$\begin{aligned} F(\rho_2^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) &\geq |\mathrm{tr}(P_2^n Q_1^n \rho_2^{\otimes n})|^2 && (7.73)\\ &\geq |\mathrm{tr}(P_2^n \rho_2^{\otimes n}) - \mathrm{tr}(P_2^n P_1^n \rho_2^{\otimes n})|^2 && (7.74)\\ &\geq |1 - \epsilon - 2^{-n(S(\rho_1)-S(\rho_2)+2\epsilon)}|^2, && (7.75) \end{aligned}$$

from which we deduce that this compression scheme is asymptotically reliable as $n \to \infty$.

We have demonstrated a compression scheme which compresses both ρ_1 and ρ_2. Notice that the measurement result r identifies, with high probability, whether the input density operator was ρ_1 or ρ_2.

More generally, suppose $\rho_1, \ldots, \rho_M$ is a set of density operators which is *entropy distinct*, that is, the density operators have distinct Von Neumann entropies. We will show that there is a compression scheme which is optimal with respect to all of these sources, by a straightforward generalization of the previous construction. In particular, order the density operators such that

$$S(\rho_1) < S(\rho_2) < \ldots < S(\rho_M). \tag{7.76}$$

Then perform the following procedure:

1. Starting with $i = 1$, perform the measurement defined by the projectors P_i^n and Q_i^n. Label the corresponding measurement result M_i; $M_i = 0$ if P_i^n occurred, and $M_i = 1$ if Q_i^n occurred. Repeat, incrementing i until $M_i = 0$ for some value of i, and then proceed to the next step. If $M_i \neq 0$ for all i, then set $M_{M+1} = 0$, and proceed to the next step.

2. Let r be the minimal value such that $M_r = 0$. If $r = M+1$ then give up. Otherwise, compress the typical subspace of ρ_r into $S(\rho_r)$ qubits per use of the source.

3. To decompress, if $r = M+1$, then do nothing. Otherwise, do the appropriate unitary transform on the $S(\rho_r)$ qubits per source symbol to move them back to the typical subspace of ρ_r.

Note that the dynamic fidelity of this compression scheme if the source is ρ_i is lower bounded by

$$F(\rho_i^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) \geq |\mathrm{tr}(P_i^n Q_{i-1}^n \dots Q_1^n \rho_i^{\otimes n})|^2. \tag{7.77}$$

Note, however, that for any set $1 \leq i_1 \leq \ldots \leq i_k \leq i-1$,

$$\left|\rho_i^{\otimes n} P_{i_1}^n \dots P_{i_k}^n P_i^n\right| \leq 2^{-n(S(\rho_i)-\epsilon)} P_i^n, \tag{7.78}$$

and therefore, for sufficiently large n and small ϵ,

$$F(\rho_i^{\otimes n}, \mathcal{D}^n \circ \mathcal{C}^n) \geq |1 - \epsilon - 2^i \epsilon|^2 \tag{7.79}$$

$$\geq \quad |1 - (2^m + 1)\epsilon|^2. \tag{7.80}$$

Letting $\epsilon \to 0$ we obtain the result.

We have shown how to compress a finite set of entropy-distinct sources. What about more general sources? In section 7.5.1 the following surprising result is proved: there exists a countable set $\pi = \{\rho_1, \rho_2, \ldots\}$ of source density operators which is both entropy distinct, and dense in the set of all source density operators.

Using previous results we can construct a compression scheme which is optimal with respect to all the source density operators in π. This is not universal data compression, but it is data compression which works for a set of sources dense in the set of all possible sources. Recall that a set X is said to be dense in Y if every point in Y can be arbitrarily well approximated by a point in X.

Fix m. Let $(\mathcal{C}_m^n, \mathcal{D}_m^n)$ be a compression scheme which is optimal with respect to the first m elements of π, namely $\rho_1, \ldots, \rho_m$. We will use this sequence of compression schemes to construct a compression scheme which works on every element in π. The essential elements of the construction are illustrated in figure 7.3.

For each m, let $n(m)$ be such that the compression scheme $(\mathcal{C}_m^n, \mathcal{D}_m^n)$ gives fidelities greater than $1 - 1/m$ for all $n \geq n(m)$. (It is clear that $n(m)$ can be chosen in such a way that it increases with m). Now for i in the range $n(m)$ to $n(m+1) - 1$, define

$$\mathcal{C}^i \equiv \mathcal{C}_m^i \tag{7.81}$$

$$\mathcal{D}^i \equiv \mathcal{D}_m^i. \tag{7.82}$$

This defines a compression scheme $(\mathcal{C}^i, \mathcal{D}^i)$ which is optimal for all source density operators in π, by construction. It is an interesting open problem to determine whether or not asymptotically reliable compression is possible for all i.i.d. sources.

7.5.1 A dense subset of density operators with distinct entropies

This subsection proves a technical theorem which is needed in our work on universal data compression. As such, it is rather ancillary to the main line of thought in the Chapter, however I find the result to be interesting and surprising in its own right.

Block length	$\{\rho_1\}$	$\{\rho_1, \rho_2\}$	$\{\rho_1, \rho_2, \rho_3\}$	...
1	$\underline{(\mathcal{C}_1^1, \mathcal{D}_1^1)}$	$(\mathcal{C}_2^1, \mathcal{D}_2^1)$	$(\mathcal{C}_3^1, \mathcal{D}_3^1)$	...
2	$\underline{(\mathcal{C}_1^2, \mathcal{D}_1^2)}$	$(\mathcal{C}_2^2, \mathcal{D}_2^2)$	$(\mathcal{C}_3^2, \mathcal{D}_3^2)$	...
3	$\underline{(\mathcal{C}_1^3, \mathcal{D}_1^3)}$	$(\mathcal{C}_2^3, \mathcal{D}_2^3)$	$(\mathcal{C}_3^3, \mathcal{D}_3^3)$	...
...	$\underline{(\ldots, \ldots)}$	$(\ldots, \ldots)$	$(\ldots, \ldots)$	...
$n(2)$	$(\mathcal{C}_1^1, \mathcal{D}_1^1)$	$\underline{(\mathcal{C}_2^1, \mathcal{D}_2^1)}$	$(\mathcal{C}_3^1, \mathcal{D}_3^1)$	...
$n(2)+1$	$(\mathcal{C}_1^1, \mathcal{D}_1^1)$	$\underline{(\mathcal{C}_2^1, \mathcal{D}_2^1)}$	$(\mathcal{C}_3^1, \mathcal{D}_3^1)$	...
...	$(\ldots, \ldots)$	$\underline{(\ldots, \ldots)}$	$(\ldots, \ldots)$	...
$n(3)$	$(\mathcal{C}_1^1, \mathcal{D}_1^1)$	$(\mathcal{C}_2^1, \mathcal{D}_2^1)$	$\underline{(\mathcal{C}_3^1, \mathcal{D}_3^1)}$	...
$n(3)+1$	$(\mathcal{C}_1^1, \mathcal{D}_1^1)$	$(\mathcal{C}_2^1, \mathcal{D}_2^1)$	$\underline{(\mathcal{C}_3^1, \mathcal{D}_3^1)}$	...
...	...	...	...	...

Figure 7.3: Construction of a data compression scheme which works on all elements in π. The vertical columns indicate a compression scheme that is optimal with respect to the set of source density operators at the head of the column. For example, the second column contains a compression scheme which is optimal with respect to both ρ_1 and ρ_2. From this table the underlined elements are used to construct a new compression scheme which is optimal with respect to all elements of π.

We begin by defining some terms. A collection of density operators is *entropy distinct* if no two density operators in the collection have the same entropy. A collection Λ of density operators is said to be *dense* in the set of all density operators if for any $\epsilon > 0$ and density operator ρ, there exists $\rho_\Lambda \in \Lambda$ such that $D(\rho, \rho_\Lambda) < \epsilon$, where $D(\cdot, \cdot)$ is the absolute distance studied in Chapter 5.

Theorem 23 *There exists a countable set $\{\rho_1, \rho_2, \ldots\}$ of density operators which is both entropy distinct, and dense in the set of all density operators.*

We will begin the proof of the theorem with some definitions and a lemma. Recall from subsection 5.2.1 that the entropy function is continuous with respect to the absolute distance. Moreover, the set of density operators is easily seen to be compact with respect to the absolute distance.

Suppose $\epsilon > 0$. Then an *ϵ-net* on the set of density operators is a collection Λ_ϵ such that given any density operator ρ, there exists $\rho_\Lambda \in \Lambda$ such that

$$D(\rho_\Lambda, \rho) < \epsilon. \tag{7.83}$$

Lemma 2 *Let $\epsilon > 0$ be given. Then there exists a finite entropy distinct ϵ-net $\{\rho_1, \ldots, \rho_n\}$.*

Proof

Since the set of density operators is compact, there exists a finite $\epsilon/2$-net $\{\sigma_1, \ldots, \sigma_n\}$. Suppose two or more of these density operators have the same entropy. For example, suppose $S(\sigma_1) = S(\sigma_2)$. Then we set $\rho_1 \equiv \sigma_1$ and perturb σ_2 by a small amount, for example

$$\rho_2 \equiv p\sigma_2 + (1-p)\frac{I}{d} \tag{7.84}$$

where $p \approx 1$ is chosen in such a way that $S(\rho_2) \neq S(\sigma_2)$, but $D(\sigma_2, \rho_2) < \epsilon/2$. (Note that if $\sigma_2 = I/d$ then it may be necessary to perturb σ_2 using a different state, say a pure state). It is easy to see that by perturbing all the density operators σ_i by a distance less than $\epsilon/2$ it is possible to ensure that the resulting perturbed set ρ_i is entropy distinct. Moreover, the resulting set is an ϵ-net, since

$$D(\rho, \rho_i) \leq D(\rho, \sigma_i) + D(\sigma_i, \rho_i) < \frac{\epsilon}{2} + \frac{\epsilon}{2}, \tag{7.85}$$

where i has been chosen such that $D(\rho, \sigma_i) < \epsilon/2$. This completes the proof of the lemma.

QED

Proof (**Main theorem**)

Returning to the main theorem, for each $n = 1, 2, \ldots$, let Λ_n be a finite entropy distinct $1/(2n)$-net. We will use these nets to construct an increasing sequence Λ'_n of $1/n$-nets. Set $\Lambda'_1 \equiv \Lambda_1$. Given Λ'_n we construct Λ'_{n+1} as follows. By perturbing each element in Λ_{n+1} by a distance at most $1/[2(n+1)]$ to create a perturbed set Λ^p_{n+1} it is possible to ensure that the resulting union

$$\Lambda'_{n+1} \equiv \Lambda'_n \cup \Lambda^p_{n+1} \tag{7.86}$$

is entropy distinct. Observing that Λ^p_{n+1} is a $1/(n+1)$-net we see that Λ'_n is a finite, entropy distinct $1/n$-net such that

$$\Lambda'_1 \subseteq \Lambda'_2 \subseteq \Lambda'_3 \subseteq \ldots. \tag{7.87}$$

Define

$$\Lambda \equiv \bigcup_{i=1}^{\infty} \Lambda'_i. \tag{7.88}$$

It is now easy to see that Λ is a countable, entropy distinct subset of the set of density operators. Furthermore, Λ is dense, since given ρ and $\epsilon > 0$, choose n such that $\frac{1}{n} < \epsilon$, and then $\rho' \in \Lambda'_N \subseteq \Lambda$ such that

$$d(\rho, \rho') < \frac{1}{n} < \epsilon. \tag{7.89}$$

This completes the proof of the theorem.

QED

Notice, incidentally, that there was very little about the entropy function that was used in the proof. All one needs is a function which varies sufficiently near any point in the space of interest in order for the result to hold. It is not even necessary that the function be continuous. Similarly, it is obviously possible to extend the result well beyond the choice of the absolute distance and the set of density operators. We will not investigate such generalizations here.

7.6 Conclusion

This Chapter has presented simplified proofs of the quantum data compression theorem, and studied universal quantum data compression. How might we extend the work further? There are several natural questions we might ask:

- What is largest class of sources for which an analogue of the typical subspace theorem holds? Any such source will necessarily satisfy the quantum data compression theorem.
- Is reliable data compression possible for all i.i.d. sources?
- Can universal data compression (or similar schemes, such as the compression scheme described in this Chapter) be implemented efficiently using a quantum circuit?

The ultimate utility of quantum data compression depends upon whether large scale quantum computers are ever built. At present, this eventuality appears to be quite far off. Nevertheless, I am hopeful that such large scale devices will one day be built, and that it may even be found useful to implement data compression in those devices. In any case, it is certainly true that studies of quantum data compression give insights into quantum information that allow us to make progress in other, perhaps more immediately practical, areas of quantum information theory.

Summary of Chapter 7: Quantum data compression

- **Quantum data compression:** We have shown that compression can be performed in such a way that the entanglement of the source with another system can be recovered with high fidelity.This is in contrast to Schumacher's original theorem, which used a weaker measure of fidelity, the ensemble fidelity, as a measure of reliability.
- **A classical side channel does not decrease the minimal resources required for reliable storage of quantum information.**
- **Universal quantum compression scheme is possible for a set of sources dense in the set of all (i.i.d.) quantum sources.**

Chapter 8

Entanglement

What is entanglement? This is a question we've skirted, until now. We've seen that entanglement is a useful resource which can be used to assist in the performance of quantum information processing tasks. It would be useful to have a more precise way of quantifying what we mean by entanglement, and understanding what it can be used to do. The purpose of this Chapter is to begin to develop such an understanding.

The most familiar example of an entangled system which we've met is the maximally entangled state of a two qubit system,

$$|\psi\rangle = \frac{1}{\sqrt{2}} \left(|00\rangle + |11\rangle \right). \tag{8.1}$$

As we have seen in earlier chapters, this state has many remarkable properties which do not appear to have classical analogues. In this Chapter we will explore in greater detail what it means for two quantum systems to be entangled.

A great deal of work has been done studying the properties of entanglement, and I will not attempt to list all of this work here. It has been widely suggested (see, for example, [93]) that entanglement is of central importance in giving quantum information processing systems an edge over classical information processing systems, although other authors [174] have argued that entanglement is at most part of the story in explaining this difference in computational power. Nevertheless, despite extensive work entanglement remains a poorly understood phenomenon.

I regret to say that I have not succeeded in understanding entanglement as deeply as I would have liked. It seems to be a difficult subject; this is reflected also in the relatively slow progress that has been made in the literature, despite much ingenuity on the part of many researchers. This Chapter is mainly concerned with reviewing a simple quantitative tool which has been developed to study entanglement, the *entanglement of formation*, and proving several new properties of this tool. The most significant new result is a relationship between the entanglement of formation and negative conditional quantum entropies. In Chapter 10 this result will be used to help gain insight into the quantum channel capacity.

The Chapter is structured as follows. In section 8.1 we study entanglement for pure states of a two part composite system. Section 8.2 develops many elementary properties of entanglement for arbitrary quantum states, including mixed states. Section 8.3 steps back from the abstract properties of entanglement studied in the preceding section, and examines some simple examples which may be used to build intuition about entanglement. Along the way, we note an amusing fact about entanglement which makes clear that it must play a rather subtle role in quantum information

processing: evidence that there may be quantum computations which cannot be simulated efficiently on a classical computer, but in which no two qubits ever become entangled. Section 8.4 concludes the Chapter with an overview of some of the future directions which could be taken to the study of entanglement in quantum systems.

Section 8.1 reviews the work of other researchers on pure state entanglement. Section 8.2 is based partially upon other people's work, however the main result of the section, the entropy-entanglement inequality, is my own work. The remainder of the Chapter is largely my own original work. I would like to thank Dorit Aharonov and Bill Wootters for several fun conversations and some great papers that got me more deeply interested in the subject, and convinced me that understanding this mysterious stuff we call entanglement is one of the central problems of quantum information theory. Their thoughts and words helped motivate much of the work here.

8.1 Pure state entanglement

What does it mean for a composite system to be entangled? In the case of pure states, an answer to this question has been given by Popescu and Rohrlich [141] for the case of a two part composite system. Suppose $|AB\rangle$ is a pure state of a two part composite system, AB. Let us assume that $E(A:B)$ is some measure of the entanglement between systems A and B. Suppose we demand that the following properties hold for E:

1. E is a function of the state $|AB\rangle$ alone.

2. E is continuous with respect to the Hilbert space distance for pure states of AB.

3. Suppose there are classical parties Alice and Bob who can manipulate systems A and B, respectively. Suppose they can perform operations on their own systems, and communicate classical information. The entanglement $E(A:B)$ between them can not be increased by such operations.

4. The entanglement is additive. That is, suppose Alice and Bob jointly possess a number of systems, $(AB)_1, (AB)_2, \ldots, (AB)_n$. The systems A_i and B_i may be entangled, but we assume that the total state of the system is a product state of the pairs $(AB)_i$. Then the total entanglement between Alice and Bob is the sum of the entanglement in each subsystem, $E(A:B) = \sum_i E(A_i:B_i)$.

Popescu and Rohrlich showed that, up to an undetermined overall constant factor, the entanglement associated to a pure state is then $E(A:B) = S(A)$, where $S(A)$ is the von Neumann entropy of subsystem A. Note that $S(A) = S(B)$ as A and B are in a pure state.

The method used by Popescu and Rohrlich may be briefly described as follows. Bennett *et al* [20, 22] had earlier considered the problem of *formation* for an entangled state. Specifically, suppose Alice and Bob want to create n copies of a pure state $|AB\rangle$, using only local operations and classical communication. The *entanglement of formation* for $|AB\rangle$ is defined to be the maximal number, c, such that if Alice and Bob are provided with $\lfloor cn \rfloor$ Bell states, then as $n \rightarrow \infty$ they can create those n copies of $|AB\rangle$ with asymptotically good fidelity. Bennett *et al* show that the entanglement of formation is equal to $S(A)$.

There is a converse process to formation, which is the *distillation* of Bell states. Suppose Alice and Bob are supplied with n copies of the state $|AB\rangle$. The *entanglement of distillation* for $|AB\rangle$ is defined to be the maximal number, c, such that as $n \rightarrow \infty$, Alice and Bob can produce $\lfloor cn \rfloor$ Bell states with asymptotically good fidelity, using local operations and classical communication.

Bennett *et al* show that the entanglement of distillation for the pure state $|AB\rangle$ is equal to the entanglement of formation, $S(A)$.

Popescu and Rohrlich's argument is to suppose we have n copies of the state $|AB\rangle$. This is then transformed by local operations and classical communication into $\lfloor nS(A) \rfloor$ approximate Bell states. Let B_e be the entanglement of a Bell state. Then we see from continuity and additivity that $nE(A:B) \geq nS(A)B_e$ must hold approximately. But we may transform the $\lfloor nS(A) \rfloor$ Bell states back into n copies of $|AB\rangle$, again by local operations and classical communication, so $nS(A)B_e \geq nE(A:B)$ must also be approximately true. Letting $n \rightarrow \infty$ we see that $E(A:B) = S(A)B_e$ holds exactly. But B_e is a constant which does not depend on the state $|AB\rangle$, so we see that, up to an overall proportionality factor, the entanglement of a pure state is uniquely determined by the above axioms to be given by $E(A:B) = S(A)$.

The obvious reaction at this point is to think "That's terrific!" After all, we understand the von Neumann entropy pretty well, so it seems as though we are well on our way to understanding entanglement in general.

Unfortunately, though, it seems to be fairly difficult to even *define* entanglement in a more general context than two-part composite systems which are in pure states. Plausible operational motivations for definitions of entanglement in more general scenarios are not difficult to generate, however researchers have had limited success in calculating with these definitions. In the next section we will study some of the best developed tools for understanding the entanglement present in an arbitrary state of two quantum systems.

8.2 Mixed state entanglement

The *entanglement of formation* [22] of a state, ρ, of a composite system AB, is defined to be the minimum number of Bell states which must be shared between A and B if they are to be able to form the state ρ using only local operations and classical communication.

More precisely, suppose we have a family of protocols, one protocol for each positive integer n, such that in the nth protocol, parties A and B start out sharing one half each of $\lfloor cn \rfloor$ Bell states, for some $c \geq 0$, and the protocols only involve local operations on systems A and B, and classical communication between A and B. The family of protocols is said to be a good entanglement forming protocol for the state ρ if the nth protocol produces the state $\rho^{\otimes n}$ with asymptotic fidelity approaching one as n approaches infinity.

The entanglement of formation has been extensively studied by Wootters and collaborators [22, 82, 195]. Modulo a problem to be discussed below, they have shown that the entanglement of formation between systems A and B is given by the expression

$$\mathcal{F}(A:B) = \min \sum_x p_x S(A_x), \tag{8.2}$$

where $\{p_x, AB_x\}$ is an ensemble of pure states generating the state AB. The minimum in the definition of the entanglement of formation is over all pure state ensembles generating the state AB.

The possible problem with equation (8.2) is whether the quantity appearing within it is *additive*[1]. Recall that the operational definition of the entanglement of formation was an asymptotic definition, expressed in terms of the creation of a large number of copies of the state AB. Strictly speaking, what is shown in [22] is that

$$\mathcal{F}_o(A:B) = \limsup_{n \rightarrow \infty} \frac{\mathcal{F}(A_1 \ldots A_n : B_1 \ldots B_n)}{n}, \tag{8.3}$$

[1]This problem was pointed out by Sandu Popescu.

where $\mathcal{F}_o$ is the entanglement of formation, operationally defined, the quantity $\mathcal{F}$ appearing on the right hand side is as defined in equation (8.2), and the system $A_1 \dots A_n B_1 \dots B_n$ is a tensor product of n copies of AB. The expression $\mathcal{F}_o(A:B)$ would be equal to $\mathcal{F}(A:B)$ if it could be shown that the quantity $\mathcal{F}(A:B)$ defined by equation (8.2) is *additive* in the sense that

$$\mathcal{F}(A_1 \dots A_n : B_1 \dots B_n) = \mathcal{F}(A_1 : B_1) + \dots + \mathcal{F}(A_n : B_n). \tag{8.4}$$

So we have two definitions of the entanglement of formation: an operational definition, based upon the number of Bell states it takes to form the state in question, and an explicit formula, equation (8.2). It is believed but not yet known that these two definitions are the same. Let $\mathcal{F}_o(A:B)$ denote the operational definition of entanglement, and $\mathcal{F}(A:B)$, the definition based upon the formula (8.2). Then

$$\mathcal{F}_o(A:B) = \limsup_{n\to\infty} \frac{\mathcal{F}(A_1, \dots, A_n : B_1, \dots, B_n)}{n}. \tag{8.5}$$

From subadditivity of the entropy and equation (8.2) it is clear that

$$\mathcal{F}_o(A:B) \leq \limsup_{n\to\infty} \frac{\mathcal{F}(A_1:B_1) + \dots + \mathcal{F}(A_n:B_n)}{n} = \mathcal{F}(A:B). \tag{8.6}$$

Note that

$$\begin{aligned} \mathcal{F}_o(A:B) &= \limsup_{n\to\infty} \frac{\min \sum_x p_x S((A_1 \dots A_n)_x)}{n} && (8.7)\\ &= \limsup_{n\to\infty} \frac{\min -\sum_x p_x S((A_1 \dots A_n)_x | (B_1 \dots B_n)_x)}{n}, && (8.8) \end{aligned}$$

where the minimum is taken over all ensembles $\{p_x, (A_1 \dots A_n B_1 \dots B_n)_x\}$ generating $A_1 \dots A_n B_1 \dots B_n$. From the subadditivity of the conditional entropy, proved in section (4.4), we see that

$$\mathcal{F}_o(A:B) \geq \min -\sum_x p_x S(A_x|B_x), \tag{8.9}$$

where now the minimum is taken over all ensembles $\{p_x AB_x\}$ generating AB. From the concavity of the conditional entropy, it follows that

$$\mathcal{F}_o(A:B) \geq -S(A|B). \tag{8.10}$$

By concavity of the entropy we also have $\mathcal{F}(A:B) \leq S(B)$. Combining this with the previous equation gives

$$\mathcal{F}(A:B) - S(A,B) \leq \mathcal{F}_o(A:B) \leq \mathcal{F}(A:B). \tag{8.11}$$

Thus for states which are nearly pure, the equation (8.2) is guaranteed to be pretty close to the operational definition for the entanglement of formation.

In [195] it is stated that numerical tests provide evidence for the conjecture that the expression given in equation (8.2) is additive, and thus is the correct operational formula for the entanglement of formation. From now on we will assume that (8.2) is indeed the correct formula for the (operational) entanglement of formation.

The entanglement has many simple, useful properties:

Theorem 24 ***(Elementary properties of entanglement)***

1. **Symmetry:**

$$\mathcal{F}(A:B) = \mathcal{F}(B:A). \tag{8.12}$$

2. **The entanglement is entropy-bounded:**

$$\mathcal{F}(A:B) \leq \min(S(A), S(B)). \tag{8.13}$$

3. **The entropy-entanglement inequality:**

$$\mathcal{F}(A:B) \geq -S(A|B); \quad \mathcal{F}(A:B) \geq -S(B|A). \tag{8.14}$$

4. **More systems means more entanglement:**

$$\mathcal{F}(A:B) \leq \mathcal{F}(A:B,C). \tag{8.15}$$

5. **Adding uncorrelated systems does not change the entanglement:**
Suppose the system AB is in a product state with system C. Then

$$\mathcal{F}(A:B,C) = \mathcal{F}(A:B). \tag{8.16}$$

Proof

The symmetry property is obvious from the definition. As already noted, the entropy-boundedness is an immediate consequence of the definition of $\mathcal{F}(A:B)$, and the concavity of the entropy. Also as noted, the entropy-entanglement inequality $\mathcal{F}(A:B) \geq -S(A|B)$ follows from the concavity of the conditional entropy.

We give two proofs that more systems means more entanglement. The first proof is from the operational definition of entanglement. If we have enough singlet pairs to make n good copies of $A:B,C$ then by throwing away all n copies of C we obtain n good copies of AB. The result follows. The result also follows easily from the entropic definition of entanglement. Suppose

$$\mathcal{F}(A:B,C) = \sum_i p_i S(\mathrm{tr}_{BC}(|\psi_i\rangle\langle\psi_i|)). \tag{8.17}$$

Define

$$\rho_i \equiv \mathrm{tr}_C(|\psi_i\rangle\langle\psi_i|), \tag{8.18}$$

and let $\{\lambda_j^i, |\psi_j^i\rangle\}$ be any ensemble for ρ_i. From the concavity of entropy it follows that

$$\begin{aligned} \sum_{ij} p_i \lambda_j^i S(\mathrm{tr}_B(|\psi_j^i\rangle\langle\psi_j^i|) &\leq \sum_i p_i S(\mathrm{tr}_B(\rho_i)) && (8.19)\\ &= \sum_i p_i S(\mathrm{tr}_{BC}(|\psi_i\rangle\langle\psi_i|)) && (8.20)\\ &= \mathcal{F}(A:B,C). && (8.21) \end{aligned}$$

This establishes the result.

To prove that adding uncorrelated systems does not change the entanglement, note first that $\mathcal{F}(A:B,C) \geq \mathcal{F}(A:B)$. Suppose $\{p_i; |AB_i\rangle\}$ is an ensemble for AB such that

$$\mathcal{F}(A:B) = \sum_i p_i S(A_i). \tag{8.22}$$

Suppose $C = \sum_j \lambda_j |j\rangle\langle j|$ is an eigenensemble decomposition for system C. Then it is clear that $\{p_i\lambda_j; |AB_i\rangle \otimes |j\rangle\}$ is an ensemble for ABC, and thus

$$\begin{aligned} \mathcal{F}(A:B,C) &\leq \sum_{ij} p_i\lambda_j S(\mathrm{tr}_{BC}(|AB_i\rangle\langle AB_i| \otimes |j\rangle\langle j|)) & (8.23)\\ &= \sum_i p_i S(\mathrm{tr}_B(|AB_i\rangle\langle AB_i|)) & (8.24)\\ &= \mathcal{F}(A:B). & (8.25) \end{aligned}$$

Thus $\mathcal{F}(A:B,C) = \mathcal{F}(A:B)$.

QED

The following simple example illustrates the importance of the entanglement for fundamental operational problems.

Suppose Alice has a composite of two systems, A and B, in her possession. She prepares the joint system AB in the state AB_x, according to some probability distribution p_x. She then gives the system B to Bob, whose task it is to determine as much information as possible about the value of x. Let I denote the maximum possible mutual information Bob can obtain about x by performing operations on B. Applying the Holevo bound, as proved in section 6.1, we see that

$$I \leq S(B) - \sum_x p_x S(B_x) = S(B) \sum_x p_x S(B_x). \tag{8.26}$$

Combining this with the observation that $\mathcal{F}(A:B) \leq \sum_x p_x S(B_x)$, we see that

$$I \leq S(B) - \mathcal{F}(A:B). \tag{8.27}$$

A similar line of reasoning shows that

$$I \leq S(A) - \mathcal{F}(A:B). \tag{8.28}$$

That is, the amount of information which Bob can obtain about the preparation of system AB is bounded by a quantity determined by the entanglement existing between those systems.

Developing such general theoretical connections between the entanglement and other quantities of practical importance is one of the great open problems in the study of entanglement. We will indicate a few more such connections in the following Chapters, but it is my hunch that many more, and deeper, connections can be found between the entanglement and other aspects of quantum information theory. It would be particularly useful to be able to connect the computational power of quantum computers, either for distributed computation, as in Chapter 6, or for straight computation, to measures of entanglement.

8.3 Entanglement: Examples

In this section we will look at some simple examples where entanglement arises naturally. These examples will give us a feel for how entanglement behaves in real physical systems. As a bonus we will obtain some clues as to how entanglement enters into quantum computation.

In general the entanglement of formation is a difficult quantity to evaluate, and no general prescription is known. However, for a bipartite system of two spins, Wootters [195] has proved a conjecture of Hill and Wootters [82], which gives an explicit prescription for evaluating the entanglement. This prescription is somewhat involved, but is straightforward in principle, and quite simple to implement on a computer. Hill and Wootters introduce the *magic basis*,

$$|a\rangle = \frac{1}{\sqrt{2}}\left(|00\rangle + |11\rangle\right) \tag{8.29}$$

$$|b\rangle = \frac{i}{\sqrt{2}}\left(|00\rangle - |11\rangle\right) \tag{8.30}$$

$$|c\rangle = \frac{i}{\sqrt{2}}\left(|01\rangle + |10\rangle\right) \tag{8.31}$$

$$|d\rangle = \frac{1}{\sqrt{2}}\left(|01\rangle - |10\rangle\right). \tag{8.32}$$

For any density operator ρ of the two spin system they define

$$R \equiv \sqrt{\sqrt{\rho}\rho^*\sqrt{\rho}}, \tag{8.33}$$

where ρ^* is the complex conjugate of ρ when ρ is expressed in the magic basis. Defining λ to be the largest eigenvalue of R, they define the *concurrence* of ρ by

$$c(\rho) \equiv \max(0, 2\lambda - \mathrm{tr}(R)). \tag{8.34}$$

The entanglement of ρ is then

$$\mathcal{F}(\rho) = H\left(\frac{1}{2} + \frac{1}{2}\sqrt{1-c^2}\right), \tag{8.35}$$

where $H(x) \equiv -x\log_2 x - (1-x)\log_2(1-x)$ is the binary Shannon entropy function.

Suppose we have an ensemble of two spin systems in equilibrium at temperature T. The state of this system is given by

$$\rho = \frac{\exp(-H/kT)}{Z}, \tag{8.36}$$

where k is Boltzmann's constant, H is the Hamiltonian for the two spin system, Z is the partition function, and we use $k = \hbar = 1$,

$$Z \equiv \mathrm{tr}(\exp(-H/T). \tag{8.37}$$

Consider a system of two spins, labeled A and B, described by the Hamiltonian

$$H = \frac{a}{2}(\sigma_z^A + \sigma_z^B) + \frac{b}{4}\vec{\sigma}_A \cdot \vec{\sigma}_B \tag{8.38}$$

a and b are real constants characterizing the internal energies of the two spins and the strength of the coupling between them, respectively.

Suppose then that we have a large number of two qubit molecules in thermodynamic equilibrium. Assuming that the intermolecular interactions are essentially negligible, so that the total system is in a product state, $\rho\otimes\ldots\otimes\rho$, it follows from the additive property of entanglement that the

total entanglement present in the system is N times the entanglement present in a single molecule, where N is the total number of molecules present in the system.

At thermal equilibrium the state of the system depends only on H/T, and this can be written in the form

$$\frac{H}{T} = \frac{\frac{1}{2}(\sigma_z^A + \sigma_z^b) + \frac{1}{4}\frac{b}{a}\vec{\sigma}_A \cdot \vec{\sigma}_B}{\frac{T}{a}}. \tag{8.39}$$

From this form we see that one of the three parameters in the problem (a, b and T) can be eliminated by using the rescaled temperature T/a and rescaled coupling strength b/a. This can be accomplished by setting $a = 1$, and using b and T, as before, which is what we do for the remainder of this section.

Written out in the magic basis the Hamiltonian of the system becomes

$$H = \begin{bmatrix} \frac{b}{4} & i & 0 & 0 \\ -i & \frac{b}{4} & 0 & 0 \\ 0 & 0 & \frac{b}{4} & 0 \\ 0 & 0 & 0 & \frac{-3b}{4} \end{bmatrix}. \tag{8.40}$$

From this form we find the energy eigenstates and corresponding eigenvalues,

$$|E_1\rangle = |d\rangle \tag{8.41}$$

$$|E_2\rangle = \frac{|a\rangle + i|b\rangle}{\sqrt{2}} = |11\rangle \tag{8.42}$$

$$|E_3\rangle = |c\rangle \tag{8.43}$$

$$|E_4\rangle = \frac{|a\rangle - i|b\rangle}{\sqrt{2}} = |00\rangle \tag{8.44}$$

$$E_1 = -\frac{3b}{4} \tag{8.45}$$

$$E_2 = \frac{b}{4} - 1 \tag{8.46}$$

$$E_3 = \frac{b}{4} \tag{8.47}$$

$$E_4 = \frac{b}{4} + 1 \tag{8.48}$$

For now we will assume that $b \geq 0$ so that $E_1 \leq E_2 \leq E_3 \leq E_4$, and return to the general case later. Given the spectrum and eigenstates it is straightforward to calculate the state of the system at temperature T,

$$\rho(T) = \frac{1}{1 + e^{b/T} + 2\cosh(1/T)} \times \begin{bmatrix} \cosh(1/T) & -i\sinh(1/T) & 0 & 0 \\ i\sinh(1/T) & \cosh(1/T) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{b/T} \end{bmatrix} \tag{8.49}$$

It is now straightforward to calculate R and to see that

$$2\lambda - \mathrm{tr}(R) = \frac{e^{b/T} - 3}{1 + e^{b/T} + 2\cosh(1/T)}, \tag{8.50}$$

from which the concurrence and entanglement follow immediately.

Note that the entanglement is non-zero for $0 < T < T_e$, where

$$T_e \equiv \frac{b}{\ln 3}. \tag{8.51}$$

Clearly the entanglement vanishes for $T \geq T_e$. We will refer to T_e as the *critical temperature*, by analogy with the physics of phase transitions – at the critical temperature, a qualitative change in the system takes place; where before, no entanglement was present in the system, now there is. Note, however, that in this example at least this change is *not* associated with the presence of long-range order in the system, as is the case in a true phase transition. The concurrence is given by the expression

$$c(b,T) = \begin{cases} \frac{e^{b/T}-3}{1+e^{b/T}+2\cosh(1/T)} & \text{if } T \leq T_e \\ 0 & \text{if } T > T_e. \end{cases} \tag{8.52}$$

Since the entanglement and the concurrence are monotonically related, our study will be focused on the concurrence, since it is easier to deal with.

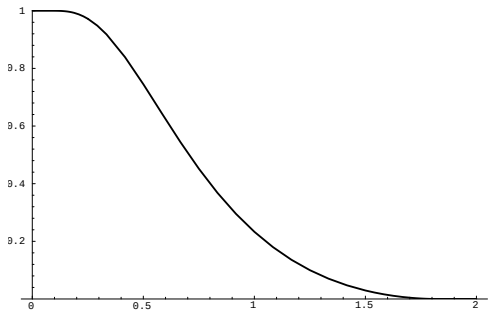

Figure 8.1: Entanglement (E) plotted as a function of temperature (T) for a coupling strength $b = 2$ in the regime where the coupling dominates.

Consider now the case where $b > 1$, which we will refer to as the *strong coupling* regime. Writing the entanglement as a function of coupling strength b and temperature T, $\mathcal{F}(b,T)$, it is clear that as $T \to 0$ we have $\mathcal{F}(b,T) \to 1$, since the ground state is the maximally entangled spin singlet

state, $|d\rangle$. Furthermore, it is easy to see that for temperatures less than the critical temperature,

$$\frac{\partial c}{\partial(1/T)} = \frac{2e^{b/T}}{(1+e^{b/T}+2\cosh(1/T))^2} \times \Big[2b + 3\sinh(1/T)e^{-b/T} + b\cosh(1/T) - \sinh(1/T)\Big]. \tag{8.53}$$

For $b \geq 1$ we have $b\cosh(1/T) \geq \cosh(1/T) > \sinh(1/T)$, from which it follows that $c(b,T)$ and thus $\mathcal{F}(b,T)$ is a *decreasing* function of temperature. It is also easy to verify that as $T \to 0$ the rate of change of $\mathcal{F}(b,T)$ with T goes to zero.

Figure 8.1 shows the entanglement plotted as a function of temperature in the strong coupling regime. As expected we see that the entanglement at $T = 0$ is 1, and then decreases monotonically to zero at T_e. For the value of the coupling strength chosen for this plot, $b = 2$, the value of the critical temperature is $T_e = 2/\ln 3 \approx 1.82$.

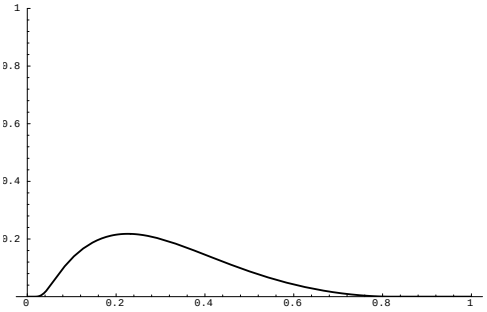

Figure 8.2: Entanglement (E) plotted as a function of temperature (T) for a weak coupling strength, $b = 0.9$.

Consider next the case where $b < 1$, the *weak coupling regime*. Once again we have

$$\frac{\partial c}{\partial(1/T)} = \frac{2e^{b/T}}{(1+e^{b/T}+2\cosh(1/T))^2} \times \Big[2b + 3\sinh(1/T)e^{-b/T} + b\cosh(1/T) - \sinh(1/T)\Big]. \tag{8.54}$$

However, it is not difficult to see by inspection that for sufficiently small T the $-\sinh(1/T)$ term in the derivative dominates, and thus $\partial c/\partial(1/T) < 0$, from which it follows that there is a regime in which the entanglement *increases* with temperature.

Figure 8.2 demonstrates this behaviour graphically. The figure shows the entanglement plotted as a function of temperature in the weak coupling regime. As expected we see that the entanglement at $T = 0$ is 0. It increases to about 0.2 near $T = 0.2$, and then decreases to zero at T_e. For the value of the coupling strength chosen for this plot, $b = 0.9$, the value of the critical temperature is $T_e = 0.9/\ln 3 \approx 0.82$.

It is not difficult to gain some intuitive feeling for why the entanglement initially increases with temperature in this case. For weak coupling the ground state of the system is the unentangled product state $|E_2\rangle = |11\rangle$. Thus, as $T \to 0$ the entanglement of the system goes to zero. For small temperatures we expect some of the population of the system to be in the first excited state, $|E_1\rangle = |d\rangle$, the maximally entangled spin singlet. The state of the system is thus largely a mixture of an unentangled state with a maximally entangled state, resulting in a small amount of entanglement.

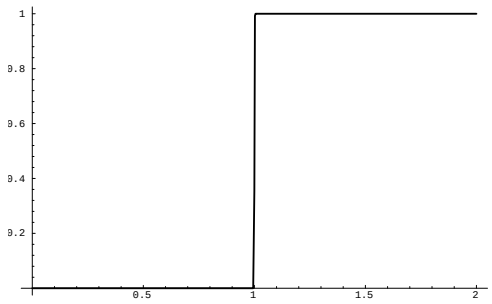

Figure 8.3: Entanglement (E) plotted as a function of the coupling strength (b) at zero temperature.

Consider now the behaviour of the entanglement as a function of the coupling strength, b. For fixed finite temperature T we see that there is a "critical value" of the coupling,

$$b_e = T \ln 3. \tag{8.55}$$

For values of the coupling above this strength the system exhibits entanglement, while for values of the coupling below this strength no entanglement exists in the system. Taking derivatives we find

$$\frac{\partial c}{\partial b} = \frac{2e^{b/T}}{T(1 + e^{b/T} + 2\cosh(1/T))^2}(2 + \cosh(1/T)), \tag{8.56}$$

which is always positive, so the entanglement *increases* as a function of b.

For $T = 0$ the preceding analysis does not apply, since the exponential terms in the denominator of the concurrence diverge as $T \to 0$. At $T = 0$ the system is in the ground state, which is the maximally entangled spin singlet state $|d\rangle$ for $b > 1$, and is the unentangled state $|11\rangle$ for $b < 1$. Thus at $T = 0$ we expect a sharp jump in the entanglement as a function of b from zero to one, at $b = 1$. Figure 8.3 shows the entanglement as a function of coupling strength at zero temperature. We clearly see such a jump in the entanglement, which is is zero below $b = 1$, and one above $b = 1$.

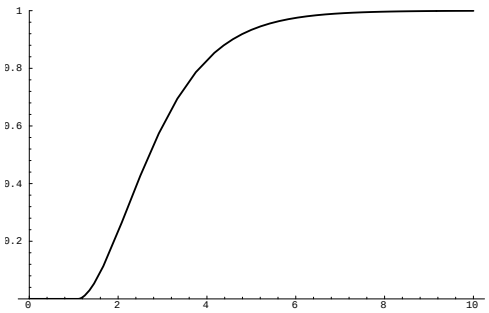

Figure 8.4: Entanglement (E) plotted as a function of the coupling strength (b) at a finite temperature, $T = 1$.

Figure 8.4 shows the entanglement as a function of coupling strength at finite temperature, in this case $T = 1$. The entanglement remains zero out to $b = 1$ at which point it suddenly begins increasing, eventually rising up to approach one for large values of the coupling.

Figure 8.5 shows the entanglement as a function of both the coupling strength and the temperature. The parameters have been chosen so that b is near the value one, where most of the interesting behaviour in this model occurs. The differing behaviour of the entanglement for $b < 1$ and $b > 1$ is clearly visible on this diagram.

Figure 8.6 shows the entanglement as a function of coupling strength and temperature, this time for very large values of the coupling and temperature. One sees plainly from this figure that as the coupling strength is increased, the critical temperature also increases.

The results in this section imply that properties associated with the entanglement of composite quantum systems are not derivable completely within the formalism of ordinary statistical mechanics. If that were the case then the properties of entanglement would be determined completely by the partition function,

$$Z \equiv \mathrm{tr}(\exp(-H/kT)). \tag{8.57}$$

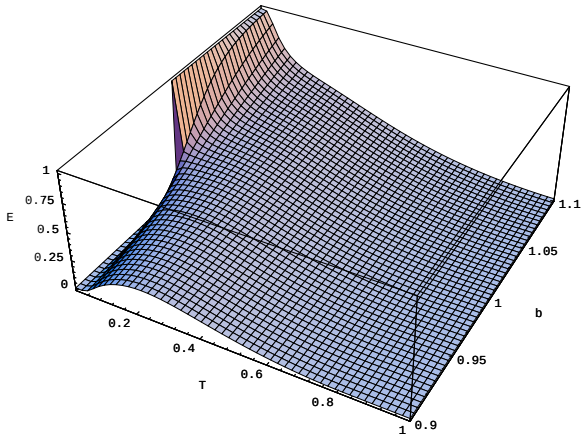

Figure 8.5: Entanglement (E) plotted as a function of temperature (T) and the coupling strength (b). This figure shows the behaviour for values of the coupling strength near the crossover at $b = 1$.

Consider a two spin system with the same energies as the spin system we have been considering, but whose energy eigenstates are unentangled product states of the two systems. It is clear that the partition functions for these two systems are the same, since they have the same energies, and thus the two systems are identical from the point of view of statistical mechanics. It is also clear that the entanglement of such a system is always zero, since its density operator is diagonal in a product basis for the system.

It seems to me that one of the major directions in which research into entanglement can be taken is to develop a theory of statistical mechanics and thermodynamics which adequately accounts for entanglement[2] in a natural fashion. It may be interesting, for example, to study transport properties of entanglement in non-equilibrium systems, or to investigate whether there is a connection between entanglement and quantum phase transitions.

A rather different direction to take the study of entanglement is in the study of the power of quantum computation. Making use of the results in this section, I will note one amusing result in this connection, associated with the behaviour of entanglement in the context of NMR quantum computing. For definiteness, we will consider the scenario defined by Knill and Laflamme [99], in

[2]Hideo Mabuchi suggested the elegant term *thermodynamics of entanglement* for such a subject to me in early 1996. We had independently been having similar thoughts about the thermodynamics and statistical mechanics of entanglement.

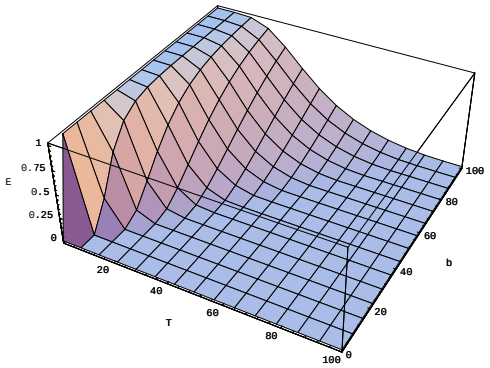

Figure 8.6: Entanglement (E) plotted as a function of temperature (T) and the coupling strength (b). This figure shows the behaviour for large values of the coupling strength, $b \gg 1$.

which they consider the power of "one bit of quantum information". Specifically, they consider a scenario in which one qubit in the state $|0\rangle$ is available, together with n qubits in the completely mixed state $I^{\otimes n}/2^n$. This can be achieved, effectively, in NMR, by making use of gradient-pulse techniques to eliminate polarization of n qubits of the molecule, and to create a pseudo-pure state on 1 qubit. We will denote the (true) initial state of the pseudo-pure qubit as ρ.

Consider a two qubit system in the completely mixed state $I^{\otimes 2}/4$. Note that $2\lambda - \mathrm{tr}(R) = -1/2$ for this state. It follows from continuity that there exists $\epsilon > 0$ such that for all states within an absolute distance ϵ of $I/4$, the entanglement of formation is zero. Note that for sufficiently high temperatures, the state $\rho \otimes I/2$ is within an absolute distance ϵ of $I^{\otimes 2}/4$. Moreover, suppose $\mathcal{E}$ is a doubly stochastic quantum operation[3] on two qubits. Then by the contractivity property of absolute distance under quantum operations and the fact that $I/4$ remains fixed under the doubly stochastic operation $\mathcal{E}$, it follows that the state $\mathcal{E}(\rho \otimes I/2)$ is also within a distance ϵ of $I/4$, and thus has zero entanglement of formation.

Next, suppose i and j are two qubits in the Knill-Laflamme scenario. Let S be the unitary operator which swaps the pseudo-pure qubit with qubit i. Note that the state of qubits i and j after performing a unitary operation U on the $n+1$ spins is given by

$$\rho'_{ij} = \mathrm{tr}_E \left(\frac{U(I^{\otimes n} \otimes \rho)U^\dagger}{2^n} \right), \tag{8.58}$$

[3]A doubly stochastic quantum operation is a complete quantum operation which preserves the completely mixed state, $\mathcal{E}(I) = I$.

where E consists of all qubits except qubits i and j. This can be rewritten as

$$\rho'_{ij} = \mathrm{tr}_E \left(\frac{US(I^{\otimes n-1} \otimes \rho_i \otimes I_j)S^\dagger U^\dagger}{2^n} \right), \tag{8.59}$$

where ρ_i indicates the state ρ, but on the ith qubit, and I_j is the identity operator on the jth qubit. It follows from the results on quantum operations in Chapter 3 that

$$\rho'_{ij} = \mathcal{E}(\rho \otimes I/2), \tag{8.60}$$

for some complete quantum operation $\mathcal{E}$. Moreover, as may be explicitly verified using equation (8.59), $I \otimes I$ is left invariant under $\mathcal{E}$, so $\mathcal{E}$ is doubly stochastic.

It follows that, in the Knill-Laflamme model of computation, with unitary operations, for sufficiently high starting temperatures, there can never be any pairwise entanglement between qubits, as measured by the entanglement of formation.

This is an interesting observation, because Knill and Laflamme [99] have found an example of simulation problems which can be solved efficiently within this model of computation which have no known efficient classical solution. Suppose no efficient classical solution is possible. Then this would be an example of a problem where quantum computers give an advantage of computational power over classical computers, yet there is never any entanglement existing between any pair of qubits. Perhaps there is entanglement in this algorithm between subsystems larger than single qubits; I do not know. However, this example does suggest that the statement "entanglement is responsible for the power of quantum computation" needs to be explored in much greater depth if it is to be made into a precise statement about the difference in computational power between quantum and classical computation. In any case, it seems as though exploration of the connection between entanglement and the power of quantum computation is any area of research that deserves considerable effort over the next few years, in order to determine to what extent entanglement is a necessary condition for quantum computers to exhibit greater computational power than classical computers.

8.4 Conclusion

In this Chapter we have discussed the elementary properties of entanglement, and made some tentative steps towards understanding its relationship with other aspects of quantum information processing. Unfortunately, I have not yet been able to integrate the subject of entanglement with other aspects of quantum information theory as smoothly as is desirable. Doing so is a very interesting challenge, which I intend to work on as part of my future research. Some questions which I intend to address include:

1. What is it about entanglement that makes it useful for quantum information processing?

2. What characteristics can be used to understand entanglement in composite systems consisting of three or more components? It is tempting to conjure up real valued measures of entanglement, just as has been done for the two-system case. I hypothesize that it may be useful to make use of more sophisticated algebraic techniques. An analogy may be helpful here. In algebraic topology [32], great progress in the study of topological spaces is made by associating algebraic constructs to topological spaces. By studying the relatively simple algebraic constructs, properties of the much more complicated topological spaces can be deduced. A similar situation may obtain in the study of entanglement.

3. How may the entanglement of formation be computed for composite systems whose components have more than two dimensions?

4. Continuous quantum phase transitions [170] occur at zero temperature in quantum systems, as some parameter in the Hamiltonian is varied. For systems with a non-degenerate ground state, the long range order associated with this phase transition must be associated with correlations arising out of entanglement in the ground state. It would be interesting to study such zero-temperature phase transitions from the point of view of quantum information theory.

5. What are the experimental signatures of entanglement? What are the simplest, most physically meaningful tests for the presence or absence of entanglement in a quantum system?

6. Can the presence of entanglement be related to other interesting physical phenomena? For example, in superconductivity, Cooper pairs form as a result of phonon exchange between electrons in a metal. For ordinary superconductors these pairs are assumed to be in the (entangled) spin singlet state. It would be interesting to know whether necessary conditions phrased in terms of entanglement can be found for the superconducting phase transition.

The study of entanglement suggests many other problems; this is only a tiny sample. I expect that this study will yield a rich and deep structure that will give us insight into both quantum information, and also into naturally occurring physical systems.

Summary of Chapter 8: Entanglement

- For pure states of AB, the entanglement is essentially unique, $E(A:B) = S(A) = S(B)$.
- **The entropy-entanglement inequality:**

$$\mathcal{F}(A:B) \geq -S(A|B).$$

- An efficient quantum algorithm with no known efficient classical analogue has been found by Knill and Laflamme [99], in which there may be no entanglement between any pair of two qubits, at any stage of the algorithm.

Chapter 9

Error correction and Maxwell's demon

Large scale quantum information processing will be enormously sensitive to the effects of noise on quantum systems. Shor [164] and Steane [172] have introduced methods for doing quantum error correction in order to preserve quantum information in the presence of noise. These methods have been developed much further by a large number of researchers, notably Gottesman [70] and Calderbank *et al* [33], who developed a powerful framework for the study of quantum codes, and by Aharonov and Ben-Or [2], Gottesman [71, 72], Kitaev [95], Knill, Laflamme and Zurek [100, 101], Preskill [144], and Shor [166], who developed methods for performing quantum information processing in the presence of noise.

In this Chapter we study quantum error correction from an information-theoretic point of view. Information-theoretic necessary and sufficient conditions for doing quantum error correction are formulated, and the information-theoretic point of view is used to study quantum error correction as a thermodynamic process, analogous to Maxwell's famous Demon [15], an information processing system which *apparently* violates the second law of thermodynamics. The material in this Chapter also serves as the basis for work in the next Chapter, on the quantum channel capacity.

Throughout the Chapter we will make heavy use of two constructions introduced earlier in this Dissertation in the study of quantum systems. First, as in Chapter 5, we assume that quantum system of interest, Q, has been purified by a second system, R, before any dynamics has occurred. The system R is assumed to undergo the trivial dynamics during any quantum process on the system Q. Moreover, quantum operations are modeled by a unitary interaction of Q with an environment, E, which is assumed to initially be in a pure state. As shown in Chapter 3, it is always possible to introduce such a model for any complete quantum operation. For incomplete quantum operations the unitary operation on QE is followed by a projection on the system E. We refer to this picture of quantum operations as the RQE picture of quantum operations.

The Chapter is structured as follows. Section 9.1 introduces the *entropy exchange*, a tool for quantifying the effects of noise in a quantum system. Section 9.2 introduces the *quantum Fano inequality*, a quantum analogue to the classical Fano inequality proved in section 6.1. Section 9.3 introduces the *coherent information*, a quantitative measure of the amount of quantum information transmitted through a quantum channel. Section 9.3 also proves the *quantum data processing inequality*, a quantum analogue of the classical data processing inequality proved in subsection 4.2.4. Section 9.4 reviews the basic concepts of quantum error correction. Section 9.5 uses the quantum Fano and data processing inequalities to obtain information-theoretic necessary and sufficient condi-

tions for quantum error correction. Section 9.6 proves several information theoretic inequalities for quantum channels. Section 9.7 formulates quantum error correction as a type of *Maxwell's demon* – a famous system proposed by Maxwell last century that was apparently able to violate the second law of thermodynamics by making observations upon a system. We do a thermodynamic analysis of quantum error correction, and show that there is no possibility of using quantum error correction to violate the second law. A consequence of our analysis, however, is that it is possible to do quantum error correction in a thermodynamically efficient manner. Section 9.8 concludes the Chapter.

Sections 9.1, 9.2 and 9.4 are largely reviews of background material. The remaining sections of the Chapter report original work, based upon collaborations with with Schumacher [158], with Caves [133], with Barnum and Schumacher [10], and with Caves, Schumacher, and Barnum [134]. Section 9.6 has not appeared elsewhere, and is an original contribution. I am especially grateful to Ben Schumacher for the many enjoyable discussions we have had about quantum information theory.

9.1 Entropy exchange

How much noise does a quantum operation cause when applied to a particular state, ρ, of a quantum system, Q? One measure of this is the extent to which the state, RQ, initially pure, becomes mixed as a result of the quantum operation. To this end, following Schumacher [157], we define the *entropy exchange* of the operation $\mathcal{E}$ upon input of ρ by

$$S_e \equiv S(\rho, \mathcal{E}) \equiv S(\rho^{RQ'}) = S(E'), \tag{9.1}$$

where the equality of the entropy exchange with $S(E')$ follows from the purity of the total state $R'Q'E'$. Thus, the entropy exchange can be regarded as the amount of entropy introduced into an initially pure environment as a result of the quantum operation $\mathcal{E}$. We use the notation S_e for the entropy exchange in situations where the arguments ρ and $\mathcal{E}$ are implied, and the notation $S(\rho, \mathcal{E})$ otherwise.

Note that the entropy exchange does not depend upon the way in which the initial state of Q, ρ, is purified into RQ. The reason is because any two purifications of Q into RQ are related by a unitary operation on the system R, as shown in appendix A. This unitary operation commutes with the action of the quantum operation on RQ, and thus the two final states of $R'Q'$ induced by the two different purifications are related by an overall unitary transformation which does not affect the entropy of $R'Q'$, giving rise to the same value for the entropy exchange. Furthermore, it follows from these results that $S(E')$ does not depend upon the particular model for $\mathcal{E}$ which is used, provided the model starts with E in a pure state.

A useful explicit formula [157] for the entropy exchange can be given, based upon the operator-sum representation for quantum operations. Suppose a complete quantum operation $\mathcal{E}$ has the operator-sum representation $\mathcal{E}(\rho) = \sum_i E_i \rho E_i^\dagger$. Then, as shown in section 3.1, a unitary model implementing this quantum operation is given by defining a unitary operator U on QE such that

$$U|\psi\rangle|0\rangle = \sum_i E_i|\psi\rangle|i\rangle, \tag{9.2}$$

where $|0\rangle$ is the initial state of the environment, and $|i\rangle$ is an orthonormal basis for the environment. Note that the state E' after application of $\mathcal{E}$ is given in this model by

$$E' = \sum_{i,j} \mathrm{tr}(E_i \rho E_j^\dagger)|i\rangle\langle j|. \tag{9.3}$$

That is, $\mathrm{tr}(E_i \rho E_j^\dagger)$ are the matrix elements of E' in the $|i\rangle$ basis. Schumacher [157] suggests defining a matrix W whose elements are given by

$$W_{ij} \equiv \mathrm{tr}(E_i \rho E_j^\dagger), \tag{9.4}$$

that is, W is the matrix of E', in an appropriate basis. This formula applies only for complete quantum operations. In the case of incomplete quantum operations a similar argument shows that the matrix elements of E' are contained in the matrix W defined by

$$W_{ij} \equiv \frac{\mathrm{tr}(E_i \rho E_j^\dagger)}{\mathrm{tr}(\mathcal{E}(\rho))}. \tag{9.5}$$

This gives rise to the useful calculational formula

$$S(\rho, \mathcal{E}) = S(W) \equiv -\mathrm{tr}(W \log W). \tag{9.6}$$

Recall from Chapter 3 that a quantum operation may have many different operator-sum representations. In particular, sets of operators E_i and F_j generate the same quantum operation if and only if $F_j = \sum_j u_{ji} E_i$, where u is a unitary matrix of complex numbers, and it may be necessary to append 0 operators to the sets E_i or F_j so that the matrix u is a square matrix.

W contains matrix elements of the environmental density operator, and thus is a positive matrix, which may be diagonalized by a unitary matrix, v, $D = vWv^\dagger$, where D is a diagonal matrix with non-negative entries. Define operators F_j by the equation

$$F_j = \sum_i v_{ji} E_i, \tag{9.7}$$

so the operators F_j give rise to the same quantum operation in the operator-sum representation. This representation of $\mathcal{E}$ gives rise to a W matrix,

$$\begin{aligned} \tilde{W}_{kl} &= \frac{\mathrm{tr}(F_k \rho F_l^\dagger)}{\mathrm{tr}(\mathcal{E}(\rho))} && (9.8) \\ &= \sum_{mn} v_{km} v_{ln}^* W_{mn} && (9.9) \\ &= D_{kl}. && (9.10) \end{aligned}$$

Thus, there is a set of operators F_j with respect to which the W matrix for the system is diagonal, with non-negative entries. Any set of operators F_j giving rise to an operator-sum representation for $\mathcal{E}$, and for which the matrix W is diagonal is said to be a *canonical representation* for $\mathcal{E}$ with respect to the input ρ. We will see later that canonical representations turn out to have a special significance for quantum error correction.

Many properties of the entropy exchange follow easily from properties of the entropy discussed in Chapter 4. For example, working in a canonical representation for a complete quantum operation, $\mathcal{E}$, on a d-dimensional space, we see immediately that $S(I/d, \mathcal{E}) = 0$ if and only if $\mathcal{E}$ is a unitary quantum operation. Therefore, $S(I/d, \mathcal{E})$ quantifies the extent to which incoherent quantum noise may occur on the system as a whole. A second example is that when $\mathcal{E}$ is restricted to be a complete quantum operation, the matrix W is easily seen to be convex-linear in ρ, and the state $R'Q'$ is convex-linear in $\mathcal{E}$. From the concavity of the von Neumann entropy it follows that $S(\rho, \mathcal{E})$ is concave in ρ and $\mathcal{E}$. Since the system RQ can always be chosen to be at most d^2 dimensional, where d is the dimension of Q, it follows that the entropy exchange is bounded above by $2 \log d$. Other properties will be derived as needed later in this Chapter, and in the next Chapter.

9.2 Quantum Fano inequality

Intuitively, if an entanglement RQ is subject to noise which results in it becoming mixed, then the fidelity of the final state $R'Q'$ with the initial state RQ cannot be perfect. Moreover, the greater the mixing, the worse the fidelity. In section 6.1 an analogous situation arose in the study of classical channels, where the uncertainty $H(X|Y)$ about the input of a channel, X, given the output, Y, was related to the probability of being able to recover the state of X from Y by the Fano inequality. Schumacher [157] has proved a very useful analogue of the classical Fano inequality, the *quantum Fano inequality*, which relates the entropy exchange and the dynamic fidelity:

$$S(\rho, \mathcal{E}) \leq h(F(\rho, \mathcal{E})) + (1 - F(\rho, \mathcal{E})) \log_2(d^2 - 1), \tag{9.11}$$

where $h(x)$ is the binary Shannon entropy. Inspection of this inequality reveals its intuitive meaning: if the entropy exchange for a process is large, then the dynamic fidelity for the process must necessarily be small, indicating that the entanglement between R and Q has not been well preserved. It will be useful to note for our later work that $0 \leq h(x) \leq 1$ and $\log(d^2 - 1) \leq 2 \log d$, so from the quantum Fano inequality,

$$S(\rho, \mathcal{E}) \leq 1 + 2(1 - F(\rho, \mathcal{E})) \log d. \tag{9.12}$$

To prove the quantum Fano inequality, consider an orthonormal set of d^2 basis states, $|\psi_i\rangle$, for the system RQ. This basis set is chosen so $|\psi_1\rangle = |RQ\rangle$. If we form the quantities $p_i \equiv \langle\psi_i|(R'Q')|\psi_i\rangle$, then from the results of subsection 4.3.3 it follows that

$$S(R'Q') \leq H(p_1, \ldots, p_{d^2}), \tag{9.13}$$

where $H(p_i)$ is the Shannon information of the set p_i. Elementary algebra shows that

$$\begin{aligned} H(p_1, \ldots, p_{d^2}) &= h(p_1) \\ &\quad + (1 - p_1) H\left(\frac{p_2}{1 - p_1}, \ldots, \frac{p_{d^2}}{1 - p_1}\right). \end{aligned} \tag{9.14}$$

Combining this with the observation that $H(\frac{p_2}{1-p_1}, \ldots, \frac{p_{d^2}}{1-p_1}) \leq \log(d^2 - 1)$ and $p_1 = F(\rho, \mathcal{E})$ by definition of the dynamic fidelity gives,

$$S(\rho, \mathcal{E}) \leq h(F(\rho, \mathcal{E})) + (1 - F(\rho, \mathcal{E})) \log(d^2 - 1), \tag{9.15}$$

which is the quantum Fano inequality.

The quantum Fano inequality has been proved using the dynamic fidelity as a measure of how well information is preserved when it is passed through a quantum channel. It is possible to give an alternative formulation of the quantum Fano inequality based upon the dynamic distance. The simplest such statement to prove is

$$S(\rho, \mathcal{E}) \leq \frac{1}{e} + D(\rho, \mathcal{E}) \log d. \tag{9.16}$$

Note that the intuitive meaning of this inequality is essentially the same as for the quantum Fano inequality based upon the dynamic fidelity: a large value for the entropy exchange implies that the dynamic distance for the process must be quite large, indicating that entanglement has not been well preserved. To prove this inequality, we make use of *Fannes' inequality*, which we proved in subsection 5.2.1. Fannes' inequality states that for two density operators A and B,

$$|S(A) - S(B)| \leq \frac{1}{e} + D(A, B) \log d. \tag{9.17}$$

Thus

$$\begin{aligned} S(\rho, \mathcal{E}) &= |S(R'Q') - S(RQ)| && (9.18) \\ &\leq \frac{1}{e} + D(R'Q', RQ) \log d && (9.19) \\ &= \frac{1}{e} + D(\rho, \mathcal{E}) \log d, && (9.20) \end{aligned}$$

which completes the proof. This inequality may easily be strengthened by making use of the stronger form of Fannes' inequality which we proved in subsection 5.2.1, at the cost of some loss in clarity.

In our work we will make use of the quantum Fano inequality based upon the dynamic fidelity, rather than the dynamic distance. Nevertheless, it is useful to keep in mind that an alternate formulation of the quantum Fano inequality is available, and may be potentially useful for some applications.

9.3 The quantum data processing inequality

In subsection 4.2.4 we reviewed a classical result about Markov processes known as the *data processing inequality*. Recall that the data processing inequality states that for a Markov process $X \rightarrow Y \rightarrow Z$,

$$H(X) \geq H(X:Y) \geq H(X:Z), \tag{9.21}$$

with equality in the first stage if and only the random variable X can be recovered from Y with probability one.

There is a quantum analogue to the data processing inequality, which Schumacher and I proved in [158]. Suppose a two stage quantum process occurs, described by quantum operations $\mathcal{E}_1$ and $\mathcal{E}_2$,

$$\rho \xrightarrow{\mathcal{E}_1} \rho' \xrightarrow{\mathcal{E}_2} \rho''. \tag{9.22}$$

We define the quantum *coherent information* by

$$I(\rho, \mathcal{E}) \equiv S(\rho') - S(\rho, \mathcal{E}). \tag{9.23}$$

This quantity, coherent information, is intended to play a role in quantum information theory analogous to the role played by the mutual information $H(X:Y)$ in classical information theory. It is not immediately apparent that the coherent information is the correct quantum analogue of the mutual information, and we will spend some time over the next two Chapters in an attempt to justify this claim.

Part of the reason for taking the coherent information seriously as a quantity like mutual information is that later in the section we prove that it satisfies the following *quantum data processing inequality*,

$$S(\rho) \geq I(\rho, \mathcal{E}_1) \geq I(\rho, \mathcal{E}_2 \circ \mathcal{E}_1), \tag{9.24}$$

with equality in the first inequality if and only if it is possible to reverse the operation $\mathcal{E}_1$, in a sense to be described below. Comparison with the classical data processing inequality shows that the coherent information plays a role in the quantum data processing inequality identical to the role played by the mutual information in the classical data processing inequality.

Such a heuristic argument can not be regarded as any sort of a rigorous justification for the view that the coherent information is the correct quantum analogue of the classical mutual information. Such a justification ought to come as a consequence of the role played by the coherent information in questions related to the quantum channel capacity. This question will be the topic of the next Chapter.

Let us return to the proof of the quantum data processing inequality. This result is proved using four systems: R, Q, E_1 and E_2. R and Q are used in their familiar roles from Chapter 5. E_1 and E_2 are systems initially in pure states, chosen such that a unitary interaction between Q and E_1 generates the dynamics $\mathcal{E}_1$, and a unitary interaction between Q and E_2 generates the dynamics $\mathcal{E}_2$. The proof of the first stage of the quantum data processing inequality is to apply the subadditivity inequality $S(R'E_1') \leq S(R') + S(E_1')$ to obtain

$$\begin{aligned} I(\rho, \mathcal{E}_1) &= S(\mathcal{E}_1(\rho)) - S(\rho, \mathcal{E}_1) && (9.25)\\ &= S(Q') - S(E_1') && (9.26)\\ &= S(R'E_1') - S(E_1') && (9.27)\\ &\leq S(R') = S(R) = S(Q) = S(\rho). && (9.28) \end{aligned}$$

The proof of the second part of the data processing inequality is to apply the strong subadditivity inequality,

$$S(R''E_1''E_2'') + S(E_1'') \leq S(R''E_1'') + S(E_1''E_2''). \tag{9.29}$$

From purity of the total state of $R''Q''E_1''E_2''$ it follows that

$$S(R''E_1''E_2'') = S(Q''). \tag{9.30}$$

Neither of the systems R or E_1 are involved in the second stage of the dynamics in which Q and E_2 interact unitarily. Thus, their state does not change during this stage: $R''E_1'' = R'E_1'$. But from the purity of RQE_1 after the first stage of the dynamics,

$$S(R''E_1'') = S(R'E_1') = S(Q'). \tag{9.31}$$

The remaining two terms in the subadditivity inequality are now recognized as entropy exchanges,

$$\begin{aligned} S(E_1'') = S(E_1') = S(\rho, \mathcal{E}_1), && (9.32)\\ S(E_1''E_2'') = S(\rho, \mathcal{E}_2 \circ \mathcal{E}_1). && (9.33) \end{aligned}$$

Making these substitutions into the inequality obtained from strong subadditivity (9.29) yields

$$S(Q'') + S(\rho, \mathcal{E}_1) \leq S(Q') + S(\rho, \mathcal{E}_2 \circ \mathcal{E}_1), \tag{9.34}$$

which can be rewritten as the second stage of the data processing inequality,

$$I(\rho, \mathcal{E}_1) \geq I(\rho, \mathcal{E}_2 \circ \mathcal{E}_1). \tag{9.35}$$

This concludes the proof of the quantum data processing inequality.

The data processing inequality will be invaluable in our study of quantum error correction, and the quantum channel capacity. To understand why it is important, consider a somewhat analogous statement, the Second Law of Thermodynamics. The constraint that the entropy of a closed system can never decrease is tremendously useful in thermodynamics. In a somewhat similar fashion,

we have obtained here a quantity (the coherent information) which is non-increasing under arbitrary quantum operations. I expect that this non-increasing property will have many uses beyond even those to which we put it in this Dissertation.

We conclude the section by noting for future reference that the first part of the data processing inequality need not hold when $\mathcal{E}_1$ is not trace-preserving. The reason for this is that it is no longer necessarily the case that $R' = R$, and thus it may not be possible to make the identification $S(R') = S(R)$. For example, suppose we have a three dimensional state space with orthonormal states $|1\rangle$, $|2\rangle$ and $|3\rangle$. Let P_{12} be the projector onto the two dimensional subspace spanned by $|1\rangle$ and $|2\rangle$, and P_3 the projector onto the subspace spanned by $|3\rangle$. Let $\rho = \frac{p}{2} P_{12} + (1-p) P_3$, where $0 < p < 1$, and $\mathcal{E}(\rho) = P_{12} \rho P_{12}$. Then by choosing p small enough we can make $S(\rho) \approx 0$, but $I(\rho, \mathcal{E}) = 1$, so we have an example of a non trace-preserving operation which does not obey the data processing inequality.

9.4 Quantum error correction

Noise is a great bane of information processing systems. Whenever possible we build our systems to avoid noise completely, and where that is not possible, we try to protect against the effects of noise. For example, components in modern computers are extremely reliable, with a failure rate typically below one error in 10^{17} operations. For most practical purposes we can act as if computer components are completely noiseless. On the other hand, many systems in widespread use do suffer from a substantial noise problem. Modems and CD players make use of error correcting codes to protect against the effects of noise. The details of the techniques used to protect against noise in practice are sometimes rather complicated, but the basic principles are easily understood. The key idea is that if we wish to protect a message against the effects of noise, then we should *encode* the message by adding some redundant information to the message. That way, even if some of the information in the encoded message is corrupted by noise, there will be enough redundancy in the encoded message that it is possible to recover or *decode* the message so that all the information in the original message is recovered.

For example, suppose we wish to send a bit from one location to another through a noisy communications channel. Suppose that the effect of the noise in the channel is to flip the bit being transmitted with probability $p > 0$; with probability $1 - p$ the bit is transmitted without error. This is known as the binary symmetric channel (see figure 9.1). A simple means of protecting the bit against the effects of noise is to replace the bit we wish to protect with three copies of itself:

$$0 \rightarrow 000 \tag{9.36}$$

$$1 \rightarrow 111. \tag{9.37}$$

We now send all three bits through the channel. At the receiver's end of the channel three bits are output, and the receiver has to decide what the value of the original bit was. Suppose 001 was output from the channel. Provided the probability p of a bit flip is not too high, it is very likely that the third bit was flipped by the channel, and that 0 was the bit that was sent.

This type of decoding is called *majority voting*, since the decoded output from the channel is whatever value, 0 or 1, appears more times in the actual channel output. Majority voting fails if two or more of the bits sent through the channel were flipped, and succeeds otherwise. The probability that two or more of the bits is flipped is $3p^2(1-p) + p^3$, so the probability of error is $p_e \equiv 3p^2 - 2p^3$. Without encoding the probability of an error was p, so the code improves matters if $p_e < p$, which occurs whenever $p < 1/2$.

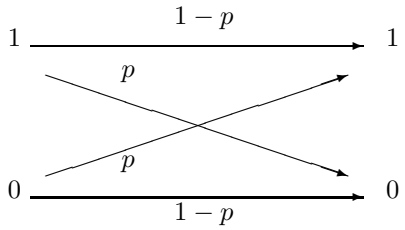

Figure 9.1: Binary symmetric channel.

The type of code we have described is called a *repetition code*, since we encode the message to be sent by repeating it a number of times. A similar technique has been used for millenia as a part of everyday conversation: if we're having difficulty understanding someone's spoken language, perhaps because they have a foreign accent, we ask them to repeat what they're saying. We may not catch all the words either time, but we can put the iterations together to produce a coherent message.

Many interesting and clever techniques have been developed in the theory of classical error correcting codes; unfortunately these techniques are beyond the scope of this Dissertation. However, the key idea is always to encode messages by adding enough redundancy that the original message is recoverable after noise has acted on the encoded message. How much redundancy needs to be added depends on how severe the noise in the channel is.

To protect quantum states against the effects of noise we would like to have *quantum error correcting codes*. This section is a review of the elementary theory of quantum error correcting codes. In the next section we will re-examine quantum error correcting codes from an information-theoretic viewpoint.

There are some important differences between classical information and quantum information that require new ideas to be introduced to make quantum error correcting codes possible:

- *No cloning*: One might try to implement the repetition code quantum mechanically by duplicating the quantum state three or more times. This is forbidden by the no cloning theorem [60, 196]. Even if cloning were possible, it would not be possible to measure and compare the three quantum states output from the channel.

- *Errors are continuous*: A continuum of different errors may occur on a single qubit. Determining which error occurred in order to correct it would appear to require infinite precision, and therefore infinite resources.

- *Measurement destroys quantum information*: In classical error correction we observe the output from the channel, and decide what decoding procedure to adopt. Observation in quantum mechanics generally destroys the quantum state under observation, and makes recovery impossible.

Suppose we send qubits through a channel which leaves the qubits untouched with probability $1 - p$, and flips the qubits with probability p. That is, with probability p the state $|\psi\rangle$ is taken to the state $X|\psi\rangle$. This channel is called the *bit flip channel*, and we will now show how to protect qubits against the effects of noise from this channel.

Suppose we encode the single qubit state $a|0\rangle + b|1\rangle$ in three qubits as $a|000\rangle + b|111\rangle$. A convenient way to write this encoding is

$$|0\rangle \rightarrow |0_L\rangle \equiv |000\rangle \tag{9.38}$$
$$|1\rangle \rightarrow |1_L\rangle \equiv |111\rangle, \tag{9.39}$$

where it is understood that superpositions of basis states are taken to corresponding superpositions of encoded states. The notation $|0_L\rangle$ and $|1_R\rangle$ indicate that these are the *logical* zero and one states, not the *physical* zero and one states.

Suppose the initial state $a|0\rangle + b|1\rangle$ has been perfectly encoded. Each of the three qubits is passed through an independent copy of the bit flip channel. Suppose a bit flip occurred on one or fewer of the qubits. There is a simple two stage *error correction* procedure which can be used to recover the correct quantum state in this case:

1. (*Error detection* or *syndrome diagnosis*). We perform a measurement which tells us what error, if any, occurred on the quantum state. The measurement result is called the *error syndrome*. For the bit flip channel there are four error syndromes, corresponding to the four projection operators

$$P_0 = |000\rangle\langle 000| + |111\rangle\langle 111| \quad \text{no error} \tag{9.40}$$
$$P_1 = |100\rangle\langle 100| + |011\rangle\langle 011| \quad \text{bit flip on qubit one} \tag{9.41}$$
$$P_2 = |010\rangle\langle 010| + |101\rangle\langle 101| \quad \text{bit flip on qubit two} \tag{9.42}$$
$$P_3 = |001\rangle\langle 001| + |110\rangle\langle 110| \quad \text{bit flip on qubit three.} \tag{9.43}$$

   Suppose, for example, that a bit flip occurs on qubit one, so the corrupted state is $a|100\rangle + b|011\rangle$. Notice that $\langle\psi|P_1|\psi\rangle = 1$ in this case, so the outcome of the measurement result (the error syndrome) is certainly 1. Notice, furthermore, that syndrome measurement does not cause any change to the state: it is $a|100\rangle + b|011\rangle$ both before and after syndrome measurement.

2. (*Recovery*). We use the value of the error syndrome to tell us what procedure can be used to recover the initial state. For example, if the error syndrome was 1, indicating a bit flip on the first qubit, then we flip that qubit again, recovering the original state $a|000\rangle + b|111\rangle$ with perfect accuracy. The four possible error syndromes and the recovery procedure in each case are: 0 (no error) – do nothing; 1 (bit flip on first qubit) – flip the first qubit again; 2 (bit flip on second qubit) – flip the second qubit again; 3 (bit flip on third qubit) – flip the third qubit again. In each case it is easy to see that the original state is recovered with perfect accuracy for each value of the error syndrome.

This error correction procedure works perfectly, provided bit flips occur on one or fewer of the three qubits. This occurs with probability $(1-p)^3 + 3p(1-p)^2 = 1 - 3p^2 + 2p^3$. The probability of an error remaining uncorrected is therefore $3p^2 - 2p^3$, just as for the repetition code we studied earlier. Once again, provided $p < 1/2$ the encoding and decoding improve the reliability of storage of the quantum state.

In some ways this error analysis is inadequate. The problem is that not all errors and states in quantum mechanics are created equal: quantum states live in a continuous space, so it is possible for some errors to corrupt a state by a tiny amount, while others mess it up completely. An extreme example is provided by the bit flip "error" X, which does not affect the state $(|0\rangle + |1\rangle)/\sqrt{2}$ at all, but flips the $|0\rangle$ state so it becomes a $|1\rangle$. In the former case we would not be worried about a bit flip error occurring, while in the latter case we would obviously be very worried.

To address this problem we make use of the *fidelity* quantity introduced in Chapter 5. Recall that the fidelity between a pure and a mixed state is given by

$$F(|\psi\rangle, \rho) \equiv \langle\psi|\rho|\psi\rangle. \tag{9.44}$$

The object of quantum error correction is to increase the fidelity with which quantum information is stored. By the results of Chapter 5, if we can perform computations with a high enough fidelity, then the measurement results output from the computation will be sufficiently close in distribution to the desired distribution to consider the computation successful.

Let's compare the *minimum* fidelity achieved by the three qubit bit flip code with the minimum fidelity achieved without error correction. Suppose the quantum state of interest is $|\psi\rangle$. Without using the error correcting code the state of the qubit after being sent through the channel is

$$\rho = (1-p)|\psi\rangle\langle\psi| + pX|\psi\rangle\langle\psi|X. \tag{9.45}$$

The fidelity is given by

$$F = \langle\psi|\rho|\psi\rangle = (1-p) + p\langle\psi|X|\psi\rangle\langle\psi|X|\psi\rangle. \tag{9.46}$$

The second term on the right hand side is non-negative. When $|\psi\rangle = |0\rangle$ the second term is zero so we see that the minimum fidelity is $F = 1-p$. Suppose the three qubit error correcting code is used to protect the state $|\psi\rangle = a|0_L\rangle + b|1_L\rangle$. The quantum state after the channel and error correction is

$$\rho = \left[(1-p)^3 + 3p(1-p)^2\right]|\psi\rangle\langle\psi| + \ldots \tag{9.47}$$

The included term is all the contributions from the correctable errors – no error at all, and a bit flip on a single qubit. The omitted terms are the contributions from bit flips on two or three qubits. The omitted terms are non-negative, so the fidelity we calculate will be a *lower bound* on the true fidelity. We see that $F = \langle\psi|\rho|\psi\rangle \geq (1-p)^3 + 3p(1-p)^2$. That is, the fidelity is at least $1-3p^2+2p^3$. We see that the fidelity of storage for the quantum state is improved provided $p < 1/2$, which is the conclusion we came to earlier based on a cruder analysis.

The bit flip code is interesting, but it does not seem to go beyond classical error correcting codes in any significant manner. A more interesting noisy quantum channel is the *phase flip* error model for a single qubit. In this error model the qubit is left alone with probability $1-p$, and the relative phase of the $|0\rangle$ and $|1\rangle$ states is flipped. More precisely, the phase flip operator Z (sometimes called the Pauli sigma z operator σ_z) is applied to the qubit with probability $p > 0$. The action of the phase flip Z is defined by $Z|0\rangle \equiv |0\rangle, Z|1\rangle \equiv -|1\rangle$. Thus the state $a|0\rangle + b|1\rangle$ is taken to the state $a|0\rangle - b|1\rangle$ under the phase flip. The reason this is called a phase flip is that the relative phase of the $|0\rangle$ and $|1\rangle$ states is flipped by the action of the phase flip operator Z.

There is no classical equivalent to the phase flip channel, since classical channels don't have any property equivalent to phase. However, there is an easy way to turn the phase flip channel into a bit flip channel. Suppose we apply the Hadamard gate immediately before and after the action of the phase flip channel. If the phase flip channel left the state alone then the additional Hadamard gates cancel out (since $H^2 = I$) and can be ignored. If the phase flip Z occurs, then the action with the Hadamard gates taken into account is $HZH = X$, which is the bit flip.

Quantum error correction for the phase flip channel can therefore be accomplished by encoding in three qubits as for the bit flip channel and then applying a Hadamard gate to each qubit to complete the encoding for the phase flip channel. The phase flip channel then acts independently on

each qubit. Finally, we error correct by applying a Hadamard gate to each qubit and then applying the usual error correction procedure for the bit flip code.

The encoded $|0\rangle$ and $|1\rangle$ for the three qubit phase flip code are thus

$$|0_L\rangle = H|0\rangle H|0\rangle H|0\rangle = |+\rangle|+\rangle|+\rangle \quad (9.48)$$

$$|1_L\rangle = H|1\rangle H|1\rangle H|1\rangle = |-\rangle|-\rangle|-\rangle, \quad (9.49)$$

where $|+\rangle \equiv H|0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and $|-\rangle \equiv H|1\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$.

Obviously this code for the phase flip channel has the same characteristics as the earlier code for the bit flip channel. In particular, the minimum fidelity for this code is the same as that for the three qubit bit flip code, and we have the same criteria for the code producing an improvement over the case with no error correction. We say that these two channels are *unitarily equivalent*, since there is a unitary operator U (in this case the Hadamard gate) such that the action of one channel is the same as the other, provided the first channel is preceded by U and followed by $U^\dagger$. These operations may be trivially incorporated into the encoding and error correction operations.

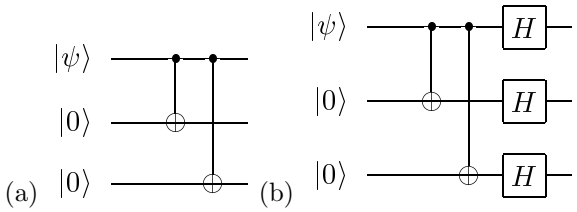

Figure 9.2: Encoding circuits for the (a) bit flip and (b) phase flip codes.

So far we've talked about encoding and error correction in the abstract. How can these operations be performed in practice? Quantum circuits for encoding the three qubit bit flip and phase flip code are shown in figure 9.2. To see that the bit flip encoding circuit works, just note that the state $|000\rangle$ is left unchanged by the circuit, while the state $|100\rangle$ is taken to the state $|111\rangle$ by the circuit. The phase flip encoding circuit is exactly the same, except we apply an extra Hadamard gate at the end of the encoding, as expected. The simplicity of design in these circuits is a general feature of many of the quantum error correcting codes which have been proposed [71], however it is by no means always the case that quantum error correction can be performed efficiently by means of a quantum circuit. One drawback of the information-theoretic approach to quantum error correction which we take later in the Chapter is that it does not seem to provide many clues about how to efficiently perform encodings and decodings.

9.4.1 Shor's code

There is a simple quantum code which can protect against the effects of *any* error, provided the error only affects a single qubit. The code is known as the *Shor code*, after its inventor [164]. The code is a combination of the three qubit phase flip and bit flip codes. We first encode the qubit using the phase flip code: $|0\rangle \rightarrow |+++\rangle, |1\rangle \rightarrow |---\rangle$. Next, we encode each of these three qubits using the bit flip code: $|+\rangle$ is encoded as $(|000\rangle + |111\rangle)\sqrt{2}$ and $|-\rangle$ is encoded as $(|000\rangle - |111\rangle)\sqrt{2}$. The result is a nine qubit code, with codewords given by:

$$|0\rangle \rightarrow |0_L\rangle \equiv \frac{(|000\rangle + |111\rangle)(|000\rangle + |111\rangle)(|000\rangle + |111\rangle)}{2\sqrt{2}} \quad (9.50)$$

$$|1\rangle \rightarrow |1_L\rangle \quad \equiv \quad \frac{(|000\rangle - |111\rangle)(|000\rangle - |111\rangle)(|000\rangle - |111\rangle).}{2\sqrt{2}}. \tag{9.51}$$

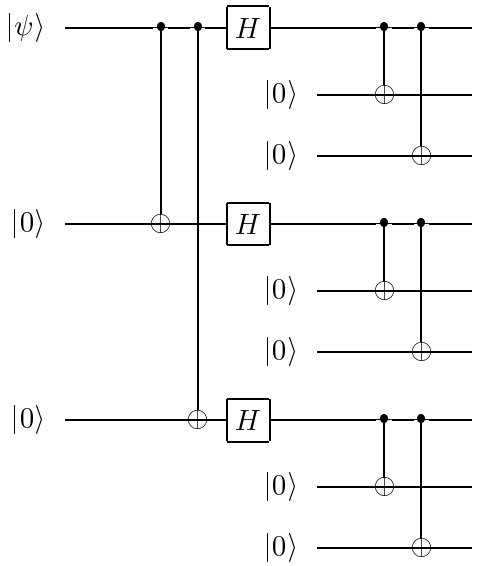

Figure 9.3: Encoding circuit for the Shor nine qubit code.

.

The quantum circuit encoding the Shor code is shown in Fig. 9.3. As described above, the first part of the circuit encodes the qubit using the three qubit phase flip code; comparison with Fig. 9.2 (b) shows that the circuits are identical. The second part of the circuit encodes each of these three qubits using the bit flip code. To do this three copies of the bit flip code encoding circuit (Fig. 9.2 (a)) is used. This method of encoding using a hierarchy of levels in this way is known as *concatenation*[97].

The Shor code is able to protect against phase flip and bit flip errors on any qubit. To see this, suppose a bit flip occurs on the first qubit. As for the bit flip code, we perform a measurement comparing the two qubits, and find that they are different. We conclude that an error occurred on the first or second qubit. As before, we do not actually measure the first and second qubit, which would destroy the coherence between them, rather, we merely compare them. Next we, compare the second and third qubit. We find that they are the same, so it could not have been the second qubit which flipped. We conclude that the first qubit must have flipped, and recover from the error by flipping the first qubit again, back to its original state.

In a similar way we can cope with a phase flip on the first qubit. We do this by comparing the *sign* of the first block of three qubits with the sign of the second block of three qubits. The phase flip on the first qubit caused the sign in the first block to be flipped, so we find that these signs are different. Next, we compare the sign of the second block of three qubits with the sign of the third block of three qubits. We find that these are the same, and conclude that the phase must have flipped in the first block of three qubits. We recover from this by flipping the sign in the first block of three qubits again, back to its original value.

Note that this procedure also allows us to recover when both a bit flip and a phase flip occur, simply by performing both procedures.

The bit and / or phase flip errors are not the only errors which the Shor code can protect against. In fact, the Shor code can protect against an *arbitrary* error, provided it only affects a single qubit. The error could even be so drastic as to remove the qubit entirely and replace it with complete garbage! The interesting thing is, no additional work needs to be done in order to protect against arbitrary errors – the procedure already described works just fine. An outline proof is as follows.

Suppose an arbitrary error occurs on the first qubit, described by a set of operators $\{E_i\}$ in some operator-sum representation (see Chapter 3). Each E_i is a single qubit operator, and thus can be expanded as a linear combination of the identity, I, the bit flip, X, the phase flip, Z, and the combined bit and phase flip, XZ:

$$E_i = e_{i0} I + e_{i1} X + e_{i2} Z + e_{i3} XZ. \tag{9.52}$$

After the noise has acted, the code is in a mixture of states, $E_i|\psi\rangle$, each of which is a superposition of the states that would have resulted if nothing had occurred (the I term in the expression for E_i), if a bit flip had occurred (the X term), if a phase flip had occurred (the Z term), or if both a bit and phase flip occurred (the XZ term). The quantum measurement used to perform error detection causes these four possible outcomes to decohere. Thus, we have a mixture of states of the form $|\psi\rangle, X|\psi\rangle, Z|\psi\rangle, XZ|\psi\rangle$. However, we have already proved that it is possible to recover the original state of the system given such a mixture, so the error correction procedure works correctly.

9.5 Information-theoretic conditions for error correction

There is an elegant set of information-theoretic conditions for quantum error correction. Suppose first that $\mathcal{E}$ is a complete quantum operation, and ρ is some input state. We will say that $\mathcal{E}$ is *perfectly reversible* upon input of ρ if there exists a complete quantum operation $\mathcal{R}$ such that

$$F(\rho, \mathcal{R} \circ \mathcal{E}) = 1. \tag{9.53}$$

From item 7 on page 103, it follows that a quantum operation is perfectly reversible if and only if for every state $|\psi\rangle$ in the support of ρ,

$$(\mathcal{R} \circ \mathcal{E})(|\psi\rangle\langle\psi|) = |\psi\rangle\langle\psi|. \tag{9.54}$$

We may connect the notion of perfect reversibility with the quantum error correcting codes which we studied in the previous section. Specifically, a quantum error correcting code was a *subspace* spanned by codewords in some larger Hilbert space. To be resilient against the noise induced by some quantum operation, $\mathcal{E}$, it is necessary that the quantum operation $\mathcal{E}$ be reversible on the subspace spanned by the codewords. Letting P be the projector onto that subspace, and d be the dimensionality, we see that the noise process $\mathcal{E}$ is correctable if and only if the operation $\mathcal{E}$ is perfectly reversible upon input of the density operator P/d.

The information-theoretic condition for a complete quantum operation $\mathcal{E}$ to be perfectly reversible upon input of ρ is that the first inequality in the quantum data processing inequality be satisfied with equality,

$$S(\rho) = I(\rho, \mathcal{E}) = S(\rho') - S(\rho, \mathcal{E}). \tag{9.55}$$

To prove necessity, suppose that $\mathcal{E}$ is perfectly reversible upon input of ρ. From the second stage of the quantum data processing inequality it follows that

$$S(\rho) - S(\rho, \mathcal{E}) \geq S(\rho'') - S(\rho, \mathcal{R} \circ \mathcal{E}). \tag{9.56}$$

From the reversibility requirement it follows that $\rho'' = \rho$. Furthermore, from the quantum Fano inequality, (9.11), and the reversibility requirement (9.53) it follows that $S(\rho, \mathcal{R} \circ \mathcal{E}) = 0$. Thus the second stage of the quantum data processing inequality may be rewritten

$$S(\rho') - S(\rho, \mathcal{E}) \geq S(\rho). \tag{9.57}$$

Combining this with the first part of the quantum data processing inequality, $S(\rho) \geq S(\rho') - S(\rho, \mathcal{E})$, we deduce that

$$S(\rho') = S(\rho) - S(\rho, \mathcal{E}), \tag{9.58}$$

for any $\mathcal{E}$ which is reversible upon input of ρ.

Next, we will give a constructive proof that satisfaction of the condition

$$S(\rho) = S(\rho') - S(\rho, \mathcal{E}) \tag{9.59}$$

implies that the quantum operation $\mathcal{E}$ is reversible upon input of ρ. Noting that $S(\rho) = S(Q) = S(R) = S(R')$, $S(\rho') = S(Q') = S(R'E')$ and $S(\rho, \mathcal{E}) = S(E')$, we see that

$$S(R') + S(E') = S(R'E'). \tag{9.60}$$

Recall from subsection 4.3.5 that this is equivalent to the condition that $R'E' = R' \otimes E'$. Suppose that the initial state of Q is $\sum_i p_i |i\rangle\langle i|$, and that we purify this state into RQ as $|RQ\rangle = \sum_i \sqrt{p_i} |i\rangle|i\rangle$. Note that $R' = R = \sum_i p_i |i\rangle\langle i|$. Furthermore, suppose that $E' = \sum_j q_j |j\rangle\langle j|$ for some orthonormal set $|j\rangle$, so that

$$R'E' = \sum_{ij} p_i q_j |i\rangle\langle i| \otimes |j\rangle\langle j|. \tag{9.61}$$

Next, we use the Schmidt decomposition to write the total state of $R'Q'E'$ after the quantum operation has been applied, as

$$|R'Q'\rangle = \sum_{ij} \sqrt{p_i q_j} |i\rangle|i,j\rangle|j\rangle, \tag{9.62}$$

where $|i,j\rangle$ is some orthonormal set of states in system Q. Define projectors P_j by

$$P_j \equiv \sum_i |i,j\rangle\langle i,j|. \tag{9.63}$$

The idea of the restoration operation is to first perform a measurement described by the projectors P_j, which reveals the state $|j\rangle$ of the environment, and then conditional on the measurement result do a unitary rotation U_j which satisfies the equation

$$U_j |i,j\rangle = |i\rangle. \tag{9.64}$$

That is, the restoration operation is given by

$$\mathcal{R}(\sigma) \equiv \sum_j U_j P_j \sigma P_j U_j^\dagger. \tag{9.65}$$

The projectors P_j are orthogonal, by the orthogonality of the states $|i,j\rangle$, but may not be complete. If this is the case, then to ensure that the quantum operation $\mathcal{R}$ is complete, it is necessary to add an extra projector $\tilde{P} \equiv I - \sum_j P_j$ to the set of projectors to make the operation complete.

Finally, note that the state of the system RQE after the reversal operation is given by

$$\begin{aligned} \sum_j U_j P_j |\psi'\rangle\langle\psi'| P_j U_j^\dagger &= \sum_j \sum_{i_1 i_2} \sqrt{p_{i_1} p_{i_2}} q_j |i_1\rangle\langle i_2| \otimes (U_j |\psi_{i_1 j}\rangle\langle\psi_{i_2 j}| U_j^\dagger) \otimes |j\rangle\langle j| & (9.66) \\ &= \sum_j \sqrt{p_{i_1} p_{i_2}} |i_1\rangle\langle i_2| \otimes |i_1\rangle\langle i_2| \otimes E', & (9.67) \end{aligned}$$

from which we see that $R''Q'' = RQ$, and thus $F(\rho, \mathcal{R} \circ \mathcal{E}) = 1$, that is, the operation $\mathcal{E}$ is perfectly reversible upon input of the state ρ, as we desired to show.

This completes the proof of the information-theoretic reversibility conditions for complete quantum operations. Some intuition about the result may be obtained by imagining that Q is a memory element in a quantum computer, R is the remainder of the quantum computer, and E is an environment whose interaction with Q causes noise. The information-theoretic reversibility condition may be stated as follows: the state of the environment, E', after the interaction, should not be correlated with the state of the remainder of the quantum computer, R', after the interaction between Q and E. That is, the environment does not learn anything about the rest of the quantum computer through interacting with Q.

We have discussed information-theoretic conditions for the perfect reversibility of complete quantum operations. It is possible to give a similar characterization of perfect reversibility for incomplete quantum operations which generalize these conditions. What does it mean for an incomplete quantum operation $\mathcal{E}$ to be perfectly reversible? As before, we will take the criterion for reversibility to be the requirement that there exist a complete quantum operation $\mathcal{R}$ such that

$$F(\rho, \mathcal{R} \circ \mathcal{E}) = 1. \tag{9.68}$$

As for the case of complete quantum operations, it is not difficult to show that this is equivalent to the condition that

$$\frac{(\mathcal{R} \circ \mathcal{E})(|\psi\rangle\langle\psi|)}{\mathrm{tr}((\mathcal{R} \circ \mathcal{E})(|\psi\rangle\langle\psi|))} = |\psi\rangle\langle\psi|, \tag{9.69}$$

for all $|\psi\rangle$ in the support of ρ.

Necessary and sufficient conditions for an arbitrary quantum operation $\mathcal{E}$ to be perfectly reversible are as follows: [133, 134]

1. There exists a constant $c > 0$ such that for all states $|\psi\rangle$ in the support of ρ, $\mathrm{tr}(\mathcal{E}(|\psi\rangle\langle\psi|)) = c$.
2. $S(\rho) = S(\rho') - S(\rho, \mathcal{E})$.

The first requirement can be given an elegant information-theoretic interpretation. We saw in Chapter 3 that incomplete quantum operations are associated with measurements on quantum systems. Suppose we think of $\mathcal{E}$ as a possible outcome that can occur as the result of a measurement on Q. Then the first requirement is just the condition that this measurement result occurs with the same

probability, regardless of which state in the support of ρ is prepared. Because of this uniformity, it follows that no information about the identity of the state is revealed to the observer through the measurement. The second requirement can be interpreted exactly as before.

To see necessity of condition 1, note that for all states σ whose support is contained within the support of ρ,

$$(\mathcal{R} \circ \mathcal{E})(\sigma) = \mathrm{tr}(\mathcal{E}(\sigma))\sigma. \tag{9.70}$$

The linearity of the left hand side implies that the right hand side must also be linear, and therefore $\mathrm{tr}(\mathcal{E}(\sigma)) = c$, for some constant c, for all states whose support is contained within the support of ρ.

To see the necessity of condition 2, we make use of the already proved necessity of condition 1. Let c be the constant value of $\mathrm{tr}(\mathcal{E}(\rho))$ from condition 1. Let E_i be a set of operators giving an operator-sum representation for $\mathcal{E}$. Let P be a projection onto the support of ρ, and $Q \equiv I - P$ the projection onto the orthocomplement of the support, that is, the kernel of ρ. Define

$$\tilde{\mathcal{E}}(\sigma) \equiv \frac{E_i P \sigma P E_i^\dagger}{c} + Q\sigma Q. \tag{9.71}$$

Note that $\tilde{\mathcal{E}}$ is a complete quantum operation such that $\tilde{\mathcal{E}}(\sigma) = \mathcal{E}(\sigma)/\mathrm{tr}(\mathcal{E}(\sigma))$ for all states σ such that the support of σ lies within the support of ρ. That is, the action of $\tilde{\mathcal{E}}$ within the support of ρ is identical to the action of $\mathcal{E}$. It follows that reversibility of $\mathcal{E}$ upon input of ρ is equivalent to the reversibility of $\tilde{\mathcal{E}}$ upon input of ρ, and therefore

$$S(\rho) = S(\rho') - S(\rho, \tilde{\mathcal{E}}). \tag{9.72}$$

But $S(\rho, \tilde{\mathcal{E}}) = S(\rho, \mathcal{E})$, since $\tilde{\mathcal{E}}$ and $\mathcal{E}$ act in the same way within the support of ρ, from which we conclude

$$S(\rho) = S(\rho') - S(\rho, \mathcal{E}), \tag{9.73}$$

as required.

To prove that these two conditions are sufficient for reversibility, construct the quantum operation $\tilde{\mathcal{E}}$ as above. As before, $\tilde{\mathcal{E}}$ and $\mathcal{E}$ have identical actions upon states in the support of ρ, and therefore $S(\rho) = S(\rho') - S(\rho, \tilde{\mathcal{E}})$. But $\tilde{\mathcal{E}}$ is a complete quantum operation, and therefore there exists a reversing operation $\mathcal{R}$ for $\tilde{\mathcal{E}}$, as constructed earlier. Since $\tilde{\mathcal{E}}$ and $\mathcal{E}$ have the same action on states in the support of ρ, it follows that $\mathcal{R}$ is a reversing operation for $\mathcal{E}$ as well. This completes the proof of the information-theoretic conditions for reversibility of a quantum operation.

Verifying the information-theoretic conditions for a specific quantum error-correcting code may not be completely trivial. It can certainly be done; for example, for the Shor code presented in the last section, however, it is usually much easier to verify that a given quantum error-correcting code works by algebraic techniques [22, 98, 71, 33]. The true benefit of the information-theoretic approach to quantum error-correction lies in the insight it gives into other problems, such as the thermodynamics of quantum error-correction, to be discussed later in this Chapter, and the quantum channel capacity, to be discussed in the next Chapter.

9.6 Information-theoretic inequalities for quantum processes

The data processing inequality is an interesting inequality with a practical use: the study of quantum error correcting codes. However, it is possible to prove many other related inequalities. In this

section I will tabulate *all* the inequalities that can be proved for a two-part quantum process, using subadditivity and strong subadditivity.

First, consider a quantum process with a single stage, described by a complete quantum operation $\mathcal{E}$.

$$\rho \stackrel{\mathcal{E}}{\longrightarrow} \rho' \tag{9.74}$$

Three systems are involved in this process, R, Q and E. Applying all possible permutations of the subadditivity inequality to $R'Q'E'$ yields three non-trivial inequalities:

$$\begin{aligned} S(R'Q') &\leq S(R') + S(Q') && (9.75)\\ S(R'E') &\leq S(R') + S(E') && (9.76)\\ S(Q'E') &\leq S(Q') + S(E'). && (9.77) \end{aligned}$$

In terms of system quantities alone these inequalities are easily rewritten

$$\begin{aligned} S(\rho, \mathcal{E}) &\leq S(\rho) + S(\rho') && (9.78)\\ S(\rho') &\leq S(\rho) + S(\rho, \mathcal{E}). && (9.79)\\ S(\rho) &\leq S(\rho') + S(\rho, \mathcal{E}). && (9.80) \end{aligned}$$

The first inequality puts an upper bound on the entropy exchange in terms of the input and output entropies. The second inequality is familiar as the first stage of the data processing inequality, slightly rewritten. The third inequality can be rewritten in a form that will be especially useful in the study of the thermodynamics of quantum error-correction,

$$\Delta S + S(\rho, \mathcal{E}) \geq 0, \tag{9.81}$$

where $\Delta S \equiv S(\rho') - S(\rho)$ is the difference between output and input entropies for the process. What this inequality tells us is that the total entropy change associated to the process is positive, where both the system Q and the environment E are included in the entropic accounting.

It is easily checked that applying all possible permutations of subadditivity and strong subadditivity to the joint system $R'Q'E'$ yields no further inequalities. Note that in all three cases the equality conditions for saturation of the inequality are obvious from the usual equality conditions for subadditivity. In particular, equality holds in (9.81) if and only if

$$Q'E' = Q' \otimes E' \tag{9.82}$$

These equality conditions are useful in our later analysis of thermodynamically efficient error correction.

Consider next the case of a two stage quantum process,

$$\rho \stackrel{\mathcal{E}_1}{\longrightarrow} \rho' \stackrel{\mathcal{E}_2}{\longrightarrow} \rho''. \tag{9.83}$$

This process involves four systems, R, Q, E_1 and E_2, where E_1 and E_2 are the environments associated with the complete quantum operations $\mathcal{E}_1$ and $\mathcal{E}_2$, respectively. Consider the state $R''Q''E_1''E_2''$. We wish to apply all possible permutations of the subadditivity and strong subadditivity inequalities to this state.

When we do this, we discover an interesting fact. The entropies of all possible subsystems of $R''Q'', E_1''E_2''$ can be expressed in terms of things we already understand, such as entropy exchanges,

with one exception. That exception is the quantity $S(Q''E_1'') = S(R''E_2'')$, which we will refer to as the *correlation entropy*, since it is a measure of the correlation existing between the environment causing noise in the first part of the dynamics, E_1, and the final state of the quantum system,

$$C(\rho, \mathcal{E}_1, \mathcal{E}_2) \equiv S(Q''E_1''). \tag{9.84}$$

To calculate the value of the correlation entropy, let $\{E_i^1\}$ and $\{E_i^2\}$ be sets of operators generating operator-sum representations for $\mathcal{E}_1$ and $\mathcal{E}_2$, respectively. Applying the usual unitary model for quantum operations, we see that the final state of $Q''E_1''E_2''$ is

$$\sum_{i_1 i_2 j_1 j_2} E_{i_2}^2 E_{i_1}^1 \rho (E_{j_1}^1)^\dagger (E_{j_2}^2)^\dagger \otimes |i_1\rangle\langle j_1| \otimes |i_2\rangle\langle j_2|. \tag{9.85}$$

Introducing an orthonormal basis $|k\rangle$ for system Q, we see that the matrix elements of $Q''E_1''$ are given by the matrix U defined by

$$U_{ik,jl} \sum_m \langle k|E_m^2 E_i^1 \rho (E_i^1)^\dagger (E_m^2)^\dagger |l\rangle, \tag{9.86}$$

and therefore $C(\rho, \mathcal{E}_1, \mathcal{E}_2) = S(U)$. We will not have any occasion to calculate correlation entropies, however it is useful to know that such an explicit formula exists which could be used to calculate such quantities if the need arises. Note also that the matrix U picks up a normalization factor of $1/\mathrm{tr}((\mathcal{E}_2 \circ \mathcal{E}_1)(\rho))$ in the case where $\mathcal{E}_2$ or $\mathcal{E}_1$ is incomplete.

Let's begin by enumerating in a table all the inequalities which can be obtained from subadditivity. To keep track of which inequalities we have evaluated, we write $(X : Y)$, where X and Y are (different) subsystems of $R''Q''E_1''E_2''$. Alongside these entries we write the corresponding entropy inequality, $S(X, Y) \leq S(X) + S(Y)$, in terms of appropriate system quantities:

Entropy Inequalities: Subadditivity

$(R'' : Q'')$	$S(\rho, \mathcal{E}_2 \circ \mathcal{E}_1)$	$\leq$	$S(\rho) + S(\rho'')$
$(R'' : E_1'')$	$S(\rho')$	$\leq$	$S(\rho) + S(\rho, \mathcal{E}_1)$
$(R'' : E_2'')$	$C(\rho, \mathcal{E}_1, \mathcal{E}_2)$	$\leq$	$S(\rho) + S(\rho', \mathcal{E}_2)$
$(Q'' : E_1'')$	$C(\rho, \mathcal{E}_1, \mathcal{E}_2)$	$\leq$	$S(\rho'') + S(\rho, \mathcal{E}_1)$
$(Q'' : E_2'')$	$S(\rho')$	$\leq$	$S(\rho'') + S(\rho', \mathcal{E}_2)$
$(E_1'' : E_2'')$	$S(\rho, \mathcal{E}_2 \circ \mathcal{E}_1)$	$\leq$	$S(\rho, \mathcal{E}_1) + S(\rho', \mathcal{E}_2)$
$(R''Q'' : E_1'')$	$S(\rho', \mathcal{E}_2)$	$\leq$	$S(\rho, \mathcal{E}_2 \circ \mathcal{E}_1) + S(\rho, \mathcal{E}_1)$
$(R''Q'' : E_2'')$	$S(\rho, \mathcal{E}_1)$	$\leq$	$S(\rho, \mathcal{E}_2 \circ \mathcal{E}_1) + S(\rho', \mathcal{E}_2)$
$(R''E_1'' : Q'')$	$S(\rho', \mathcal{E}_2)$	$\leq$	$S(\rho') + S(\rho'')$
$(R''E_1'' : E_2'')$	$S(\rho'')$	$\leq$	$S(\rho') + S(\rho', \mathcal{E}_2)$
$(R''E_2'' : Q'')$	$S(\rho, \mathcal{E}_1)$	$\leq$	$C(\rho, \mathcal{E}_1, \mathcal{E}_2) + S(\rho'')$
$(R''E_2'' : E_1'')$	$S(\rho'')$	$\leq$	$C(\rho, \mathcal{E}_1, \mathcal{E}_2) + S(\rho, \mathcal{E}_1)$
$(Q''E_1'' : R'')$	$S(\rho', \mathcal{E}_2)$	$\leq$	$C(\rho, \mathcal{E}_1, \mathcal{E}_2) + S(\rho)$
$(Q''E_1'' : E_2'')$	$S(\rho)$	$\leq$	$C(\rho, \mathcal{E}_1, \mathcal{E}_2) + S(\rho', \mathcal{E}_2)$
$(Q''E_2'' : R'')$	$S(\rho, \mathcal{E}_1)$	$\leq$	$S(\rho') + S(\rho)$
$(Q''E_2'' : E_1'')$	$S(\rho)$	$\leq$	$S(\rho') + S(\rho, \mathcal{E}_1)$
$(E_1''E_2'' : R'')$	$S(\rho'')$	$\leq$	$S(\rho, \mathcal{E}_2 \circ \mathcal{E}_1) + S(\rho)$
$(E_1''E_2'' : Q'')$	$S(\rho)$	$\leq$	$S(\rho, \mathcal{E}_2 \circ \mathcal{E}_1) + S(\rho'')$

Next, we construct a table containing all the inequalities obtainable directly from the strong subadditivity inequality, $S(X,Y,Z)+S(Y) \leq S(X,Y)+S(Y,Z)$. At the start of each row we write $(X:Y:Z)$ to indicate to which three subsystems of $R''Q''E_1''E_2''$ the strong subadditivity inequality is being applied:

Entropy Inequalities: Strong Subadditivity

$(R'':Q'':E_1'')$	$S(\rho',\mathcal{E}_2)+S(\rho'')$	$\leq$	$S(\rho,\mathcal{E}_2\circ\mathcal{E}_1)+C(\rho,\mathcal{E}_1,\mathcal{E}_2)$
$(Q'':E_1'':R'')$	$S(\rho',\mathcal{E}_2)+S(\rho,\mathcal{E}_1)$	$\leq$	$C(\rho,\mathcal{E}_1,\mathcal{E}_2)+S(\rho')$
$(E_1'':R'':Q'')$	$S(\rho',\mathcal{E}_2)+S(\rho)$	$\leq$	$S(\rho'')+S(\rho,\mathcal{E}_2\circ\mathcal{E}_1)$
$(R'':Q'':E_2'')$	$S(\rho,\mathcal{E}_1)+S(\rho'')$	$\leq$	$S(\rho,\mathcal{E}_2\circ\mathcal{E}_1)+S(\rho')$
$(Q'':E_2'':R'')$	$S(\rho,\mathcal{E}_1)+S(\rho',\mathcal{E}_2)$	$\leq$	$S(\rho')+C(\rho,\mathcal{E}_1,\mathcal{E}_2)$
$(E_2'':R'':Q'')$	$S(\rho,\mathcal{E}_1)+S(\rho)$	$\leq$	$C(\rho,\mathcal{E}_1,\mathcal{E}_2)+S(\rho,\mathcal{E}_2\circ\mathcal{E}_1)$
$(R'':E_1'':E_2'')$	$S(\rho'')+S(\rho,\mathcal{E}_1)$	$\leq$	$S(\rho')+S(\rho,\mathcal{E}_2\circ\mathcal{E}_1)$
$(E_1'':E_2'':R'')$	$S(\rho'')+S(\rho',\mathcal{E}_2)$	$\leq$	$S(\rho,\mathcal{E}_2\circ\mathcal{E}_1)+C(\rho,\mathcal{E}_1,\mathcal{E}_2)$
$(E_2'':R'':E_1'')$	$S(\rho'')+S(\rho)$	$\leq$	$C(\rho,\mathcal{E}_1,\mathcal{E}_2)+S(\rho')$
$(Q'':E_1'':E_2'')$	$S(\rho)+S(\rho,\mathcal{E}_1)$	$\leq$	$C(\rho,\mathcal{E}_1,\mathcal{E}_2)+S(\rho,\mathcal{E}_2\circ\mathcal{E}_1)$
$(E_1'':E_2'':Q'')$	$S(\rho)+S(\rho',\mathcal{E}_2)$	$\leq$	$S(\rho,\mathcal{E}_2\circ\mathcal{E}_1)+S(\rho')$
$(E_2'':Q'':E_1'')$	$S(\rho)+S(\rho'')$	$\leq$	$S(\rho')+C(\rho,\mathcal{E}_1,\mathcal{E}_2)$

It is, perhaps, slightly unfortunate that we do not make use of this plethora of entropy inequalities. Certainly, it is interesting to peruse these tables of entropy inequalities, attempting to discern the significance of each of these results. I hope that some of them may have a role to play in future research into quantum information theory.

9.7 Quantum error correction and Maxwell's demon

Error correction may decrease the entropy of a quantum system, so it is natural to inquire about the thermodynamic efficiency of this process. In this section we discuss the question of the entropy cost of error correction and show that error correction can be regarded as a sort of *refrigeration*, wherein information about the system dynamics, obtained through measurement, is used to keep the system cool. Indeed, the method of operation of an error correction scheme is very similar to that of a famous old paradox of thermodynamics, the *Maxwell demon* paradox [15] introduced by Maxwell last century, and the methods we will use to analyze the thermodynamics of quantum error correction are based upon those used by Bennett [17, 15] to resolve the paradox.

9.7.1 Error-correction by a "Maxwell demon"

Consider the error-correction "cycle" depicted in figure 9.4. The cycle can be decomposed into four stages:

1. The system, starting in a state ρ, is subjected to a noisy quantum evolution that takes it to a state ρ^n. We denote the change in entropy of the system during this stage by ΔS. In typical scenarios for error correction, we are interested in cases where $\Delta S \geq 0$, though this is not necessary.

2. A "demon" performs a measurement on the state ρ^n. We will suppose that the measurement can be described by quantum operations $\mathcal{E}_m(\rho) = M_m \rho M_m^\dagger$. As shown in section 9.5, the error detection stage of quantum error correction can always be performed in such a way. The probability that the demon obtains result m is

$$p_m = \mathrm{tr}(M_m \rho^n M_m^\dagger) \; , \tag{9.87}$$

and the state of the system conditioned on result m is

$$\rho_m = M_m \rho^n M_m^\dagger / p_m \; . \tag{9.88}$$

3. The demon "feeds back" the result m of the measurement as a unitary operation V_m that creates a final system state

$$\rho_m^c = V_m \rho_m V_m^\dagger = V_m M_m \rho^n M_m^\dagger V_m^\dagger / p_m \; , \tag{9.89}$$

In the case of error correction this final state is the "corrected" state. The state of the system, averaged over all possible measurement outcomes, is given by

$$\rho_c \equiv \sum_m p_m \rho_m^c . \tag{9.90}$$

4. The cycle is restarted. In order that this actually be a cycle and that it be a successful error correction, we must have $\rho^c = \rho$.

The second and third stages are the "error-correction" stages. The idea of error correction is to restore the original state of the system during these stages. In this section we show that the reduction in the system entropy during the error-correction stages comes at the expense of entropy production in the environment, which is at least as large as the entropy reduction.

To investigate the balance between the entropy reduction of the system and entropy production in the environment, we adopt what Caves [35] has termed the "inside view" of the demon. The "outside view" of the demon regards it as a specific physical system. By contrast, the only aspect of the Demon relevant from the "inside view" are its properties as an information processing system; it appears to itself as a set of decohered classical bits stored in some memory. After stage 3 the only record of the measurement result m is the record in the demon's memory. To reset its memory for the next cycle, the demon must erase its record of the measurement result. Associated with this erasure is a thermodynamic cost, the *Landauer erasure cost* [106], which corresponds to an entropy increase in the environment. The erasure cost of information is equivalent to the thermodynamic cost of entropy, when entropy and information are measured in the same units, conveniently chosen to be bits. Bennett [17] used the idea of an erasure cost to resolve the paradox of Maxwell demons, and Zurek [204] and later Caves [34] showed that a correct entropic accounting from the "inside view" can be obtained by quantifying the amount of information in a measurement record by the algorithmic information content I_m of the record. Algorithmic information is the information content of the most compressed form of the record, quantified as the length of the shortest program that can be used to generate the record on a universal computer. We show here that the average thermodynamic cost of the demon's measurement record is at least as great as the entropy reduction achieved by error correction.

In a particular error-correction cycle where the demon obtains measurement result i, the total thermodynamic cost of the error-correction stages is $I_m + \Delta S^c$, where

$$\Delta S^c \equiv S(\rho^c) - S(\rho^n) \tag{9.91}$$

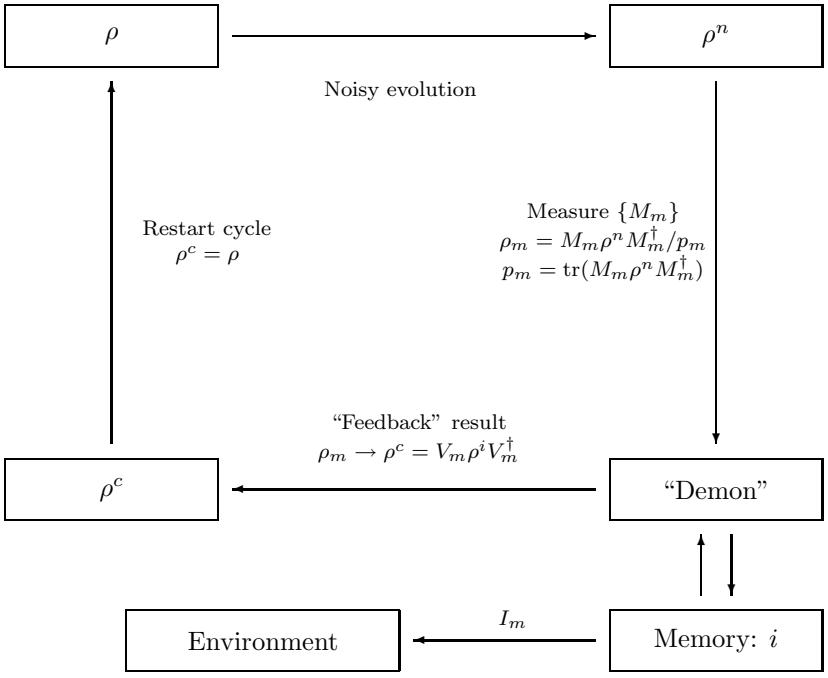

Figure 9.4: Error correction cycle.

is the change in the system entropy in the error-correction stages. Note that the stage change in the quantum system results in a change in entropy of ΔS^c, not $S(\rho_m^c) - S(\rho^n)$, because the result of the measurement record is erased after the error-correction stage, leaving the quantum system in the state $\rho^c \equiv \sum_m p_m \rho_m^c$. What is of interest to us is the average thermodynamic cost,

$$\sum_m p_m I_m + \Delta S^c \ , \tag{9.92}$$

where the average is taken over the probabilities for the measurement results. To bound this average thermodynamic cost, we now proceed through a chain of three inequalities.

The first inequality is a strict consequence of algorithmic information theory: the average algorithmic information of the measurement records is not less than the Shannon information for the probabilities p_m, that is,

$$\sum_m p_m I_m \geq H(p_m) = -\sum_m p_m \log p_m \ . \tag{9.93}$$

Furthermore, Schack [151] has shown that any universal computer can be modified to make a new universal computer that has programs for all the raw measurement records which are at most one bit longer than optimal code words for the measurement records. On such a modified universal computer, the average algorithmic information for the measurement records is within one bit of the Shannon information H.

To obtain the second and third inequalities, notice that the corrected state ρ^c can be written as

$$\rho^c = \sum_m p_m V_m \rho_m V_m^\dagger = \sum_m V_m M_m \rho^n M_m^\dagger V_m^\dagger \equiv \mathcal{R}(\rho^n) \ , \tag{9.94}$$

where $\mathcal{R}$ is the deterministic reversal operation for the error-correction stages. The operators $V_m M_m$ make up an operator-sum decomposition for the reversal operation. The probabilities p_m are the diagonal elements of the W matrix for this decomposition,

$$p_m = \mathrm{tr}(M_m \rho^n M_m^\dagger) = \mathrm{tr}(V_m M_m \rho^n M_m^\dagger V_m^\dagger) \ . \tag{9.95}$$

From the results of subsection 4.3.3 we see that

$$H(p_m) \geq S(\rho^n, \mathcal{R}), \tag{9.96}$$

with equality if and only if the operators $V_m M_m$ are a canonical decomposition of $\mathcal{R}$ with respect to ρ^n. We stress that different measurements and conditional unitaries at stages 2 and 3 lead to the same reversal operation, but may yield quite different amounts of Shannon information.

The third inequality is obtained by applying the inequality (9.81) to $\mathcal{R}$ and ρ^n:

$$S(\rho^n, \mathcal{R}) + \Delta S^c \geq 0 \ . \tag{9.97}$$

This inequality is automatically satisfied with equality if $\mathcal{R}$ error corrects $\mathcal{E}$. To see this, recall the data processing inequality gives $S(\rho) \geq S(\rho^n) - S(\rho, \mathcal{E}) \geq S(\rho^c) - S(\rho, \mathcal{R} \circ \mathcal{E})$. But the error correcting property implies that these inequality hold with equality, $\rho^c = \rho$, and $S(\rho, \mathcal{R} \circ \mathcal{E}) = 0$. Therefore we have $S(\rho^n) - S(\rho, \mathcal{E}) = S(\rho^c)$. But $S(\rho, \mathcal{R} \circ \mathcal{E}) = 0$, from which we deduce that $S(\rho, \mathcal{E}) = S(\rho^n, \mathcal{R})$, and therefore $SS(\rho^n, \mathcal{R}) + \Delta S^c = 0$ for error-correction.

Combining the three inequalities, we see that the total entropy produced during the error-correction process is greater than or equal to zero:

$$\sum_m p_m I_m + \Delta S^c \geq H(p_m) + \Delta S^c \geq S(\rho^n, \mathcal{R}) + \Delta S^c c \geq 0 \ . \tag{9.98}$$

Stated another way, this result means that the total entropy change around the cycle is at least as great as the initial change in entropy ΔS, which is caused by the first stage of the dynamics. The error-correction stage can be regarded as a kind of refrigerator, similar to a Maxwell demon, achieving a reduction in system entropy at the expense of an increase in the entropy of the environment due to the erasure of the demon's measurement record.

How then does this error-correction demon differ from an ordinary Maxwell demon? An obvious difference is that the error-correction demon doesn't extract the work that is available in the first step of the cycle as the system entropy increases under the noisy quantum evolution. A subtler, yet more important difference lies in the ways the two demons return the system to a standard state, so that the whole process can be a cycle. For the error-correction demon, it is the error-correction steps that reset the system to a standard state, which is then acted on by the noisy quantum evolution. For an ordinary Maxwell demon, the noisy quantum evolution restores the system to a standard state, typically thermodynamic equilibrium, starting from different input states representing the different measurement outcomes.

Can this error correction be done in a thermodynamically efficient manner? Is there a strategy for error correction that achieves equality in the Second Law inequality (9.98)? The answer is yes, and we give such a strategy here. The proof of the Second Law inequality (9.98) uses three inequalities, $\sum_m p_m I_m \geq H$, $H \geq S_e$, and $S_e \geq -\Delta S$. To achieve thermodynamically efficient error correction, it is necessary and sufficient that the equality conditions in these three inequalities be achieved.

We have already noted that Schack has shown that the first inequality, $\sum_m p_m I_m \geq H(p_m)$, can be saturated to within one bit by using a universal computer that is designed to take advantage of optimal coding of the raw measurement records i. On such a universal computer the average amount of space needed to store the programs for the measurement records—that is, the encoded measurement records—is within one bit of the Shannon information H. Moreover, it is possible to reduce this one bit asymptotically to zero by the use of block coding and reversible computation. The demon stores the results of its measurements using an optimal code for a source with probabilities p_m. Thus the demon stores an encoded list of measurement results. Immediately before performing a measurement, the demon decodes the list of measurement results using reversible computation. It performs the measurement, appends the result to its list, and then re-encodes the enlarged list using optimal block coding done by reversible computation. In the asymptotic limit of large blocks, the average length of the compressed list of measurement results becomes arbitrarily close to $H(p_m)$ per measurement result.

The second inequality, $H(p_m) \geq S(\rho^n, \mathcal{R})$, can be saturated by letting the measurement operators M_m and conditional unitaries V_m be those defined by the canonical decomposition of the reversal operation $\mathcal{R}$. It should be noted that the optimal method of encoding the measurement records depends on the probabilities p_m, which in turn are ultimately determined by the initial state ρ. Thus the type of encoding needed to efficiently store the measurement record generally depends on the initial state ρ. The probabilities p_m of the measurement results cannot depend on the initial state, ρ, by the results of section 9.5. It follows that for some states with support in the coding space, this error correction scheme is not thermodynamically efficient.

The third inequality, $S(\mathcal{R}, \rho^n) \geq -\big(S(\rho^c) - S(\rho^n)\big)$, is satisfied, as we have already seen, by any error-correction procedure that corrects errors perfectly. It would be interesting to see whether equality can be achieved in inequality (9.97) by error-correction schemes that do not correct errors perfectly.

9.7.2 Discussion

Zurek [205], Milburn [126], and Lloyd [117] have analyzed examples of quantum Maxwell demons, though not in the context of error correction. Lloyd notes that "creation of new information" in a quantum measurement is an additional source of inefficiency in his scheme, which involves measuring Z for a spin in a static magnetic field applied along the z axis, in order to extract energy from it. If the spin is measured in the "wrong" basis – for example, if it is initially in a pure state not an eigenstate of Z – the measurement fails to extract all the available free energy of the spin, because of the disturbance to the system state induced by the measurement. In the case of error correction, something similar happens, but it is not disturbance to the system that is the source of the inefficiency. Instead, if the ancilla involved in the reversal decoheres in the wrong basis – that is, the measurement performed by the demon is not the one defined by the canonical decomposition of the reversal operation – then the Landauer erasure cost is greater than the efficient minimum S_e. This can be thought of as creation of new information, due to disturbance of the ancilla, but the change in the system state is independent of the basis in which the ancilla decoheres.

Error correction can be accomplished in ways other than that depicted in figure 9.4. The "inside view" of the preceding subsection, in which the demon makes a measurement described by some decomposition of the reversal operation, arises when the demon is decohered by an environment, the particular measurement being defined by the basis in which the environment decoheres. If the demon is isolated from everything except the system and is initially in a pure state, then its entropy gain is $S_e = -\Delta S$ for the error-correction process. One can restart the error-correction cycle by discarding the demon and bringing up a new demon, the result being an increase in the environment's entropy by the demon's entropy S_e. This way of performing error correction, which does not involve any measurement records, is equivalent to the "outside view" of the demon's operation.

The "inside view" of the demon's operation, we stress again, arises if the demon's memory is decohered by interaction with an environment, the measurement record thus becoming "classical information." In this case the demon has the entropy $H(p_m)$ of the measurement record, not just the entropy S_e. Once this decoherence is taken into account, the different decompositions of the reversal operation, corresponding to different measurements, constitute operationally different ways of reversing things, rather than just different interpretations of the same overall interaction. Keeping in mind the variety of decompositions of the reversal operation might lead one to consider a greater variety of experimental realizations, some of which may be easier to perform than others. As we emphasize above, a reversal in which the decohered measurement results correspond to a canonical decomposition of the reversal operation is the reversal method that is most efficient thermodynamically.

9.8 Conclusion

In this Chapter we have shown that information-theoretic tools can be a powerful tool to understand quantum noise, and quantum error correction. Information-theoretic necessary and sufficient conditions for quantum error correction have been formulated, and a thermodynamic analysis of quantum error correction performed, which shows that quantum error correction functions as a kind of "quantum Maxwell's demon", for reducing the entropy of a quantum system, through observation and feedback.

Summary of Chapter 9: Error correction and Maxwell's demon

- **Entropy exchange:** Measure of noise a quantum process induces in a state.

$$S_e \equiv S(\rho, \mathcal{E}) \equiv S(\rho^{RQ'}).$$

- **Quantum Fano inequality:** A large entropy exchange implies a low dynamic fidelity.

$$S(\rho, \mathcal{E}) \leq h(F(\rho, \mathcal{E})) + (1 - F(\rho, \mathcal{E})) \log(d^2 - 1).$$

- **Coherent information:** Quantum analogue of the mutual information.

$$I(\rho, \mathcal{E}) \equiv S(\mathcal{E}(\rho)) - S(\rho, \mathcal{E}).$$

- **Data processing inequality:**

$$S(\rho) \geq I(\rho, \mathcal{E}_1) \geq I(\rho, \mathcal{E}_2 \circ \mathcal{E}_1).$$

  Equality is satisfied in the first inequality if and only if it is possible to perfectly error correct $\mathcal{E}_1$ on the subspace supporting ρ.

- **Error correction as a Maxwell's demon:** Extracting classical information about a quantum system, we can reduce its entropy, at the cost of having to *erase* the classical measurement results. This is the thermodynamic cost of quantum error correction; there is always a way for doing quantum error correction in a thermodynamically efficient way.

Chapter 10

The quantum channel capacity

A central result of Shannon's classical theory of information [160, 162, 54] is the *noisy channel coding theorem*. This result provides an *effective procedure* for determining the *capacity* of a noisy channel - the maximum rate at which classical information can be reliably transmitted through the channel.

This Chapter has two goals. The first goal is to develop general techniques for proving upper bounds on the capacity of a noisy quantum channel, which are applied to several different classes of quantum noisy channel problems. Second, I point out some of the essentially new features that quantum mechanics introduces into the noisy channel problem, which make it more difficult than the classical noisy channel problem. It is worth emphasizing at this point that this Chapter does *not* provide an effective procedure for calculating the capacity of a quantum channel, or even for calculating bounds on the channel capacity, except in very simple cases. What it represents is progress on understanding the quantum channel capacity from the point of view of the von Neumann entropy and related tools. The Chapter is based upon work done in collaboration with Schumacher [158], and with Barnum and Schumacher [10]. As the work was being carried out, independent work on the problem was being done by Lloyd [118], Bennett *et al* [22], and Shor and Smolin [166]. Additional work done since that time will be pointed out within the Chapter. The Chapter reports original work; there is little review material in the Chapter.

The Chapter is organized as follows. In section 10.1 we give a basic introduction to the problem of the noisy quantum channel, and explain the key concepts. Section 10.2 shows how the classical noisy channel coding theorem can be put into the quantum language, and explains why the capacities that arise in this context are not directly useful for applications such as quantum computing. Section 10.3 discusses the *coherent information* introduced in the previous Chapter as an analogue to the concept of *mutual information* in classical information theory. Many new results about the coherent information are proved, and we show that quantum entanglement allows the coherent information to have properties which have no classical analogue. These properties are critical to understanding what is essentially quantum about the quantum noisy channel coding problem. Section 10.4 brings us back to noisy channel coding, and formally sets up the class of noisy channel coding problems we consider. Section 10.5 proves a variety of upper bounds on the capacity of a noisy quantum channel, depending on what class of coding schemes one is willing to allow. This is followed in section 10.6 by a discussion of the achievability of these upper bounds and of other work on channel capacity. Section 10.7 formulates the new problem of a noisy quantum channel with measurement, allowing classical information about the environment to be obtained by measurement, and then used during the decoding process. Upper bounds on the corresponding channel capacity are proved. Finally, section 10.8 concludes with a summary of our results, a discussion of the new

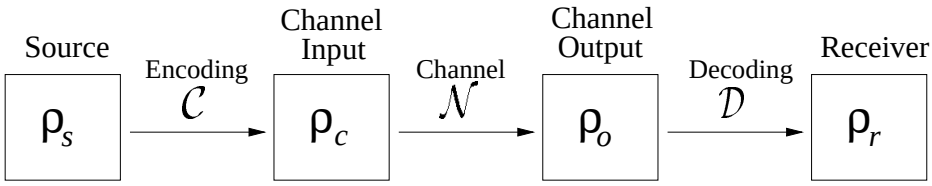

Figure 10.1: The noisy quantum channel, together with encodings and decodings.

features which quantum mechanics adds to the problem of the noisy channel, and suggestions for further research.

10.1 Noisy channel coding

The problem of noisy channel coding will be outlined in this section. Precise definitions of the concepts used will be given in later sections. The procedure is illustrated in figure 10.1.

There is a *quantum source* emitting unknown quantum states, which we wish to transmit through the channel to some receiver. Unfortunately, the channel is usually subject to noise, which prevents it from transmitting states with high fidelity. For example, an optical fiber suffers losses during transmission. Another important example of a noisy quantum channel is the memory of a quantum computer. There the idea is to transmit quantum states *in time*. The effect of transmitting a state from time t_1 to t_2 can be described as a noisy quantum channel. Quantum teleportation can also be described as a noisy quantum channel whenever there are imperfections in the teleportation process, as shown in section 3.3.

The idea of noisy channel coding is to encode the quantum state emitted by the source, ρ_s, which one wishes to transmit, using some *encoding operation*, which we denote $\mathcal{C}$. The encoded state is then sent through the channel, whose operation we denote by $\mathcal{N}$. The output state of the channel is then *decoded* using some *decoding operation*, $\mathcal{D}$. The objective is for the decoded state to match with high fidelity the state emitted by the source. As in the classical theory, we consider the fidelity of large blocks of material produced by repeated emission from the source, and allow the encoding and decoding to operate on these blocks. A channel is said to transmit a source reliably if a sequence of block-coding and block-decoding procedures can be found that approaches perfect fidelity in the limit of large block size.

Shannon's classical noisy coding theorem is proved for *discrete memoryless channels*. Discrete means that the channel only has a finite number of input and output states. By analogy we define a discrete quantum channel to be one which has a finite number of Hilbert space dimensions. In the classical case, memoryless means that the output of the channel is independent of the past, conditioned on knowing the state of the source. Quantum mechanically we take this to mean that the output of the channel is completely determined by the encoded state of the source, and is not affected by the previous history of the source.

Phrased in the language of quantum operations, we assume that there is a quantum operation, $\mathcal{N}$, describing the dynamics of the channel. The input ρ_i of the channel is related to the output ρ_o by the equation

$$\rho_i \to \rho_o = \mathcal{N}(\rho_i). \tag{10.1}$$

For the majority of this Chapter we assume, as in the previous equation, that the operation describing the action of the channel is a complete quantum operation. This corresponds to the physical assumption that no classical information about the state of the system or its environment is obtained by an external classical observer. However, in section 10.7 we go beyond this to consider the case of a noisy channel which is being observed by some classical observer, which will cause us to make use of incomplete quantum operations.

What then is the *capacity* of such a discrete memoryless quantum channel - the highest rate at which information can be reliably transmitted through the channel? The goal of a *channel capacity theorem* is to provide a procedure to answer this question. This procedure must be an *effective procedure*, that is, an explicit algorithm to evaluate the channel capacity. Such a theorem comes in two parts. One part proves an upper bound on the rate at which information can be reliably transmitted through the channel. The other part demonstrates that there are coding and decoding schemes which attain this bound, which is therefore the channel capacity. We do not prove such a channel capacity theorem in this Chapter. We do, however, derive bounds on the rate at which information can be sent through a noisy quantum channel.

Before we proceed to the more technical sections of the Chapter, it is useful to settle on a few notational conventions. Generically we denote quantum operations by $\mathcal{E}$ and the dimension of the quantum system Q by d. $\mathcal{N}$ is used to denote noisy quantum channels, which are also quantum operations. We work in the RQE picture of quantum operations, as in the previous chapter. A prime always denotes a *normalized* state. For instance,

$$R'Q' = \frac{(\mathcal{I}_R \otimes \mathcal{E})(RQ)}{\mathrm{tr}((\mathcal{I}_R \otimes \mathcal{E})(RQ))}. \tag{10.2}$$

Other notational conventions will be introduced as we proceed further.

10.2 Classical noisy channels in a quantum setting

In this section we show how classical noisy channels can be formulated in terms of quantum mechanics. We begin by reviewing the formulation in terms of classical information theory.

A classical noisy channel is described in terms of distinguishable channel states, which we label by x. If the input to the channel is symbol x then the output is symbol y with probability $p_{y|x}$. The channel is assumed to act independently on each input. For each x, the probability sum rule $\sum_y p_{y|x} = 1$ is satisfied. These *conditional probabilities* $p_{y|x}$ completely describe the classical noisy channel.

Suppose the input to the channel, x, is represented by some classical random variable, X, and the output by a random variable Y. Shannon showed that the capacity of a noisy classical channel is given by the expression

$$C_S = \max_{p(x)} H(X:Y), \tag{10.3}$$

where $H(X:Y)$ is the Shannon mutual information between X and Y, as defined in subsection 4.2.3, and the maximum is taken over all possible distributions $p(x)$ for the channel input, X. Notice that although this is not an explicit expression for the channel capacity in terms of the conditional probabilities $p_{x|y}$, the maximization can easily be performed using well known techniques from numerical mathematics. That is, Shannon's result provides an effective procedure for computing the capacity of a noisy classical channel.

All these results may be re-expressed in terms of quantum mechanics. We suppose the channel has some preferred orthonormal basis, $|x\rangle$, of signal states. For convenience we assume the

set of input states, $|x\rangle$, is the same as the set of output states, $|y\rangle$, of the channel, although more general schemes are possible. For the purpose of illustration the present level of generality suffices. A classical input random variable, X, corresponds to an input density operator for the quantum channel,

$$\rho_X \equiv \sum_x p(x)|x\rangle\langle x|. \tag{10.4}$$

The statistics of X are recoverable by measuring ρ_X in the $|x\rangle$ basis. Defining operators E_{xy} by

$$E_{xy} \equiv |y\rangle\langle x|, \tag{10.5}$$

we find that the channel operation defined by

$$\mathcal{N}(\rho) \equiv \sum_{xy} p_{y|x} E_{xy} \rho E_{xy}^\dagger. \tag{10.6}$$

is a trace-preserving quantum operation, and that

$$\mathcal{N}(\rho_X) = \rho_Y = \sum_y p(y)|y\rangle\langle y|, \tag{10.7}$$

where ρ_Y is the density operator corresponding to the random variable Y that would have been obtained from X given a classical channel with probabilities $p_{y|x}$. This gives a quantum mechanical formalism for describing classical sources and channels. It is interesting to see what form the mutual information and channel capacity take in the quantum formalism.

Notice that

$$H(X) = S(\rho_X) \tag{10.8}$$
$$H(Y) = S(\rho_Y) = S(\mathcal{N}(\rho_X)). \tag{10.9}$$

Next we compute the entropy exchange associated with the channel operating on input ρ_X, by computing the W matrix given by equation (9.5). The W matrix corresponding to the channel with input ρ_X has entries

$$W_{(xy)(x'y')} = \delta_{x,x'}\delta_{y,y'}p(x)p(y|x), \tag{10.10}$$

But the joint distribution of (X, Y) satisfies $p(x)p(y|x) = p(x, y)$. Thus W is diagonal with eigenvalues $p(x, y)$, so the entropy exchange is given by

$$S(\rho_X, \mathcal{N}) = H(X, Y). \tag{10.11}$$

It follows that

$$H(X : Y) = S(\rho_X) + S(\mathcal{N}(\rho_X)) - S(\rho_X, \mathcal{N}), \tag{10.12}$$

and thus the Shannon capacity C_S of the classical channel is given in the quantum formalism by

$$C_S = \max_{\rho_X} \left[S(\rho_X) + S(\mathcal{N}(\rho_X)) - S(\rho_X, \mathcal{N}) \right], \tag{10.13}$$

where the maximization is over all input states for the channel, ρ_X, which are diagonal in the $|x\rangle$ basis.

The problem we have been considering is that of transmitting a discrete set of orthogonal states (the states $|x\rangle$) through the channel. In many quantum applications one is not only interested in transmitting a discrete set of states, but rather the entanglement of a quantum source with another system. For this purpose we will use the *dynamic fidelity* of Chapter 5 as a figure of merit for how reliable transmission is. The capacity in this scenario is defined to be the highest rate at which quantum source can be transmitted through a noisy quantum channel, in the sense of having asymptotically high dynamic fidelity. It is easy to see, and we will show explicitly later on, that this cannot be done by considering the transmission of a set of orthogonal pure states alone. That is, the transmission of entanglement is a much more stringent condition than the transmission of classical information which we have been considering here, and consequently, the channel capacity for transmission of quantum entanglement – the main subject of this Chapter – may in general be somewhat lower than the channel capacity for transmission of classical information.

10.3 Coherent information

In this section we investigate in more detail the *coherent information*, defined in section 9.3, where it was suggested that the coherent information plays a role in quantum information theory analogous to the role played by mutual information in classical information theory; that is, suppose we consider a process defined by an input ρ, and output ρ', with the process described by a quantum operation, $\mathcal{E}$,

$$\rho \stackrel{\mathcal{E}}{\rightarrow} \rho' = \mathcal{E}(\rho). \tag{10.14}$$

I assert that the coherent information, defined by

$$I(\rho, \mathcal{E}) \equiv S\left(\frac{\mathcal{E}(\rho)}{\mathrm{tr}(\mathcal{E}(\rho))}\right) - S(\rho, \mathcal{E}), \tag{10.15}$$

plays a role in quantum information theory analogous to that played by the mutual information $H(X : Y)$ in classical information theory, where X is the input to a classical channel, and Y is the output from that channel. Heuristic arguments for why this is so were given in the previous Chapter. Of course, the true justification for regarding the coherent information as the quantum analogue of the mutual information is its success as the quantity appearing in results on channel capacity, as discussed in later sections. This is the appropriate motivation for all definitions in information theory, whether classical or quantum: their success at quantifying the resources needed to perform some interesting physical task, not some abstract mathematical motivation.

Subsection 10.3.1 studies in detail the properties of the coherent information. In particular, we prove several results related to convexity that are useful both as calculational aids, and also for proving later results. Subsection 10.3.2 proves the *entropy-fidelity* lemma that glues together many of our later proofs of upper bounds on the channel capacity. Finally, subsections 10.3.3 and 10.3.4 describe two important ways the behaviour of the coherent information differs from the behaviour of the mutual information when quantum entanglement is allowed.

10.3.1 Properties of coherent information

The set of quantum operations forms a positive cone, that is, if $\mathcal{E}_i$ is a collection of quantum operations and λ_i is a set of non-negative numbers then $\sum_i \lambda_i \mathcal{E}_i$ is also a quantum operation. In

this section we prove two very useful properties of the coherent information. First, it is easy to see that for any quantum operation $\mathcal{E}$ and non-negative λ,

$$I(\rho, \lambda\mathcal{E}) = I(\rho, \mathcal{E}). \tag{10.16}$$

This follows immediately from the definition of the coherent information. A slightly more difficult property to prove is the following.

Theorem 25 ***(convexity theorem for coherent information)***
Suppose $\mathcal{E}_i$ are quantum operations. Then

$$I(\rho, \sum_i \mathcal{E}_i) \leq \frac{\sum_i tr(\mathcal{E}_i(\rho)) I(\rho, \mathcal{E}_i)}{tr(\sum_i \mathcal{E}_i(\rho))}. \tag{10.17}$$

This result will be extremely useful in our later work. An important and immediate corollary is the following:

Corollary 2 *If a complete quantum operation, $\mathcal{E} = \sum_i p_i \mathcal{E}_i$ is a convex sum ($p_i \geq 0, \sum_i p_i = 1$) of complete quantum operations $\mathcal{E}_i$, then the coherent information is convex,*

$$I(\rho, \sum_i p_i \mathcal{E}_i) \leq \sum_i p_i I(\rho, \mathcal{E}_i). \tag{10.18}$$

The proof of the corollary is immediate from the theorem.

Proof (convexity theorem for coherent information)

The theorem follows from the concavity of the *conditional* entropy, Corollary 1, on page 84. By definition

$$I(\rho, \mathcal{E}) = S(Q') - S(R'Q') = -S(R'|Q'). \tag{10.19}$$

The theorem follows immediately from the concavity of the conditional entropy.

QED

The following lemma, from [123], is extremely useful in computing the maxima of convex functions on convex sets. Later in this Chapter we will be interested in the computation of such maxima.

Lemma 3 *Suppose f is a continuous convex function on a compact, convex set, S. Then there is an extremal point at which f attains its global maximum.*

The proof is obvious. The reason for our interest in the proof is because for fixed ρ and complete quantum operations $\mathcal{E}$, the coherent information $I(\rho, \mathcal{E})$ is a convex, continuous function of the operation $\mathcal{E}$, as just shown. The set of trace-preserving quantum operations forms a compact, convex set, and thus by the convexity lemma, $I(\rho, \mathcal{E})$ attains its maximum for a quantum operation $\mathcal{E}$ which is extremal in the set of all trace-preserving quantum operations.

A further useful result concerns the additivity of coherent information,

Theorem 26 ***(additivity for independent channels)***
Suppose $\mathcal{E}_1, \ldots, \mathcal{E}_n$ are quantum operations and $\rho_1, \ldots, \rho_n$ are density operators. Then

$$I(\rho_1 \otimes \ldots \otimes \rho_n, \mathcal{E}_1 \otimes \ldots \mathcal{E}_n) = \sum_i I(\rho_i, \mathcal{E}_i). \tag{10.20}$$

The proof is immediate from the additivity property of entropies for product states.

10.3.2 The entropy-fidelity lemma

The following lemma is the glue which holds together much of our later work on proving upper bounds to channel capacities. In this section we will prove the lemma only for the special case of complete quantum operations. A similar but more complicated result is true for general quantum operations, and will be given in section 10.7.

Lemma 4 ***(entropy-fidelity lemma)***

Suppose $\mathcal{E}$ is a complete quantum operation, and ρ is some quantum state. Then for all complete quantum operations $\mathcal{D}$,

$$S(\rho) \leq I(\rho, \mathcal{E}) + 2 + 4(1 - F(\rho, \mathcal{D} \circ \mathcal{E})) \log d. \tag{10.21}$$

This lemma is extremely useful in obtaining proofs of bounds on the channel capacity. In order for the dynamic fidelity to be close to one, the quantity appearing on the right hand side must be close to zero. This shows that the entropy of ρ cannot greatly exceed the coherent information $I(\rho, \mathcal{E})$ if the dynamic fidelity of the total process – $\mathcal{E}$ followed by $\mathcal{D}$ – is to be close to one.

Proof

To prove the lemma, notice that by the second part of the data processing inequality, (9.24),

$$S(\rho) - I(\rho, \mathcal{E}) \quad \leq \quad S(\rho) - S((\mathcal{D} \circ \mathcal{E})(\rho)) + S(\rho, \mathcal{D} \circ \mathcal{E}). \tag{10.22}$$

Applying inequality (9.81) gives

$$S(\rho) - S((\mathcal{D} \circ \mathcal{E})(\rho)) \leq S(\rho, \mathcal{D} \circ \mathcal{E}), \tag{10.23}$$

and combining the previous two inequalities gives

$$\begin{aligned} S(\rho) - I(\rho, \mathcal{E}) &\leq 2S(\rho, \mathcal{D} \circ \mathcal{E}) && (10.24) \\ &\leq 2h(F(\rho, \mathcal{D} \circ \mathcal{E})) + 2(1 - F(\rho, \mathcal{D} \circ \mathcal{E})) \log(d^2 - 1), && (10.25) \end{aligned}$$

where the second step follows from the quantum Fano inequality, (9.11). But the binary Shannon entropy h is bounded above by 1 and $\log(d^2 - 1) \leq 2 \log d$, so

$$S(\rho) \leq I(\rho, \mathcal{E}) + 2 + 4(1 - F(\rho, \mathcal{D} \circ \mathcal{E})) \log d. \tag{10.26}$$

This completes the proof.

QED

The inequality in the statement of the entropy-fidelity lemma is strong enough to prove the asymptotic bounds of most interest in our later work. The somewhat stronger inequality (10.25) is also useful when proving one-shot results, that is, when no block coding is being used. We will not make any use of it in this Chapter.

10.3.3 Quantum characteristics of the coherent information I

There are at least two important respects in which the coherent information behaves differently from the classical mutual information. In this subsection and the next we will explain what these differences are.

Classically, suppose we have a Markov process,

$$X \to Y \to Z. \tag{10.27}$$

Intuitively we expect that

$$H(X:Z) \leq H(Y:Z), \tag{10.28}$$

and, indeed, in subsection 4.2.4 we proved this "data pipelining inequality", based on the definition of the mutual information. The idea is that any information about X that reaches Z must go through Y, and thus will also be information that Z has about Y. However, the quantum mechanical analogue of this result fails to hold. We shall see that the reason it fails is due to quantum entanglement.

Example 1:

Suppose we have a two-part quantum process described by quantum operations $\mathcal{E}_1$ and $\mathcal{E}_2$.

$$\rho \rightarrow \mathcal{E}_1(\rho) \rightarrow (\mathcal{E}_2 \circ \mathcal{E}_1)(\rho). \tag{10.29}$$

Then, in general

$$I(\rho, \mathcal{E}_2 \circ \mathcal{E}_1) \not\leq I(\mathcal{E}_1(\rho), \mathcal{E}_2). \tag{10.30}$$

An explicit example showing that this is the case will be given below. It is not possible to prove *any* general inequality of this sort for the coherent information - examples may be found where a $<, >$ or $=$ sign could occur in the last equation. We will now show how the purely quantum mechanical effect of entanglement is responsible for this property of coherent information.

Observe first that the truth of the equation

$$I(\rho, \mathcal{E}_2 \circ \mathcal{E}_1) \leq I(\mathcal{E}_1(\rho), \mathcal{E}_2), \tag{10.31}$$

is equivalent to

$$S(\mathcal{E}_1(\rho), \mathcal{E}_2) \leq S(\rho, \mathcal{E}_2 \circ \mathcal{E}_1). \tag{10.32}$$

This last equation makes it easy to see why (10.31) may fail. It is because the entropy of the joint environment for processes $\mathcal{E}_1$ and $\mathcal{E}_2$ (the quantity on the right-hand side) may be less than the entropy of the environment for process $\mathcal{E}_2$ alone (the quantity on the left). This is a property peculiar to quantum mechanics, which is caused by entanglement; there is no classical analogue. In particular, the entropy-entanglement inequality on page 151 showed that the entanglement between E_1'' and E_2'' satisfies

$$\mathcal{F}(E_1'' : E_2'') \geq S(E_2'') - S(E_1'', E_2'') = S(\mathcal{E}_1(\rho), \mathcal{E}_2) - S(\rho, \mathcal{E}_2 \circ \mathcal{E}), \tag{10.33}$$

demonstrating that entanglement between E_1'' and E_2'' must exist in order that (10.31) be violated.

An explicit example where this is the case will now be given. For convenience we will do so in the language of coding and channel operations, since this is the language that will be most convenient later. $\mathcal{E}_1$ is to be identified with the coding operation, $\mathcal{C}$, and $\mathcal{E}_2$ is to be identified with the channel operation, $\mathcal{N}$.

Suppose we have a four dimensional state space. We will suppose we have an orthonormal basis $|1\rangle, |2\rangle, |3\rangle, |4\rangle$, and that P_{12} is the projector onto the space spanned by $|1\rangle$ and $|2\rangle$, and P_{34} is the projector onto the space spanned by $|3\rangle$ and $|4\rangle$. Let U be a unitary operator defined by

$$U \equiv |3\rangle\langle 1| + |4\rangle\langle 2| + |1\rangle\langle 3| + |2\rangle\langle 4|. \tag{10.34}$$

The channel operation is defined by

$$\mathcal{N}(\rho) \equiv P_{12}\rho P_{12} + U^\dagger P_{34}\rho P_{34} U, \tag{10.35}$$

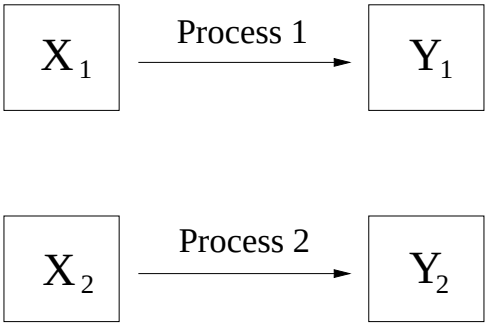

Figure 10.2: Dual classical channels operating on inputs X_1 and X_2 produce outputs Y_1 and Y_2.

and we use an encoding defined by

$$\mathcal{C}(\rho) \equiv \frac{1}{2}P_{12}\rho P_{12} + \frac{1}{2}UP_{12}\rho P_{12}U^\dagger + P_{34}\rho P_{34}. \tag{10.36}$$

It is easily checked that for any state ρ whose support lies wholly in the space spanned by $|1\rangle$ and $|2\rangle$,

$$(\mathcal{N} \circ \mathcal{C})(\rho) = \rho. \tag{10.37}$$

It follows that

$$I(\rho, \mathcal{N} \circ \mathcal{C}) = S(\rho). \tag{10.38}$$

It is also easy to verify that

$$I(\mathcal{C}(\rho), \mathcal{N}) = 2S(\rho) - 1. \tag{10.39}$$

Thus there exist states ρ such that

$$I(\rho, \mathcal{N} \circ \mathcal{C}) > I(\mathcal{C}(\rho), \mathcal{N}), \tag{10.40}$$

providing an example of (10.30).

10.3.4 Quantum characteristics of the coherent information II

The second important difference between coherent information and classical mutual information is related to the property known classically as *subadditivity of mutual information*. Suppose we have several independent channels operating. Figure 10.2 shows the case of two channels.

These channels are numbered $1, \ldots, n$ and take as inputs random variables $X_1, \ldots, X_n$. The channels might be separated spatially, as shown in the figure, or in time. The channels are assumed to act independently on their respective inputs, and produce outputs $Y_1, \ldots, Y_n$. It is not difficult to show that [54] (c.f. page 72)

$$H(X_1, \ldots, X_n : Y_1, \ldots, Y_n) \leq \sum_i H(X_i : Y_i). \tag{10.41}$$

This property is known as the *subadditivity* of mutual information. It is used, for example, in proofs of the weak converse to Shannon's noisy channel coding theorem. We will show that the corresponding quantum statement about coherent information fails to hold.

Example 2: There exists a quantum operation $\mathcal{E}$ and a density operator ρ_{12} such that

$$I(\rho_{12}, \mathcal{E} \otimes \mathcal{E}) \not\leq I(\rho_1, \mathcal{E}) + I(\rho_2, \mathcal{E}), \tag{10.42}$$

where $\rho_1 \equiv \mathrm{tr}_2(\rho_{12})$ and $\rho_2 \equiv \mathrm{tr}_1(\rho_{12})$ are the usual reduced density operators for systems 1 and 2.

An example of (10.42) is the following. Suppose system 1 consists of two qubits, A and B. System 2 consists of two more qubits, C and D. As the initial state we choose

$$\rho_{12} = \frac{I_A}{2} \otimes |BD\rangle\langle BD| \otimes \frac{I_C}{2}, \tag{10.43}$$

where $|BD\rangle$ is a Bell state shared between systems B and D.

The action of the channel on A and B is as follows: it sets bit B to some standard state, $|0\rangle$, and allows A through unchanged. This is achieved by swapping the state of B out into the environment. Formally,

$$\mathcal{E}(\rho_{AB}) = \rho_A \otimes |0\rangle\langle 0|. \tag{10.44}$$

The same channel is now set to act on systems C and D:

$$\mathcal{E}(\rho_{CD}) = \rho_C \otimes |0\rangle\langle 0|. \tag{10.45}$$

A straightforward though slightly tedious calculation shows that with this channel setup

$$I(\rho_1, \mathcal{E}) = I(\rho_2, \mathcal{E}) = 0, \tag{10.46}$$

and

$$I(\rho_{12}, \mathcal{E} \otimes \mathcal{E}) = 2. \tag{10.47}$$

Thus this setup provides an example of the violation of subadditivity for the coherent information, (10.42).

10.4 Noisy channel coding revisited

In this section we return to noisy channel coding. Recall the basic procedure for noisy channel coding, as illustrated in figure 10.3.

Suppose a quantum source has output ρ_s. A quantum operation, which we shall denote $\mathcal{C}$, is used to *encode* the source source, giving the input state to the channel, $\rho_i \equiv \mathcal{C}(\rho)$. The encoded state is used as input to the noisy channel, giving a channel output $\rho_o \equiv \mathcal{N}(\rho_i)$. Finally, a decoding quantum operation, $\mathcal{D}$, is used to decode the output of the channel, giving a *received state*, $\rho_r \equiv \mathcal{D}(\rho_o)$. The goal of noisy channel coding is to find out what source states can be sent with high dynamic fidelity. That is, we want to know for what states ρ_s can encoding and decoding operations be found such that

$$F(\rho_s, \mathcal{D} \circ \mathcal{N} \circ \mathcal{C}) \approx 1. \tag{10.48}$$

Typically, it is the entropy of a state which determines whether it can be sent with high dynamic fidelity. If large blocks of source states with entropy R per use of the source can be sent through

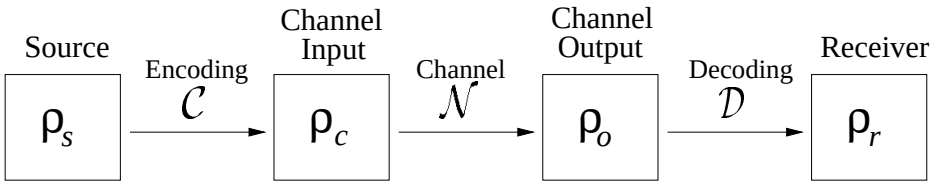

Figure 10.3: The noisy quantum channel, together with encodings and decodings.

equally large blocks of channel with high dynamic fidelity, we say the channel is *transmitting at the rate R*.

Shannon's noisy channel coding theorem is an example of a *channel capacity* theorem. Such theorems come in two parts:

1. First an *upper bound* is placed on the rate at which information can be sent reliably through the channel. This upper bound should be expressible entirely in terms of channel quantities.

2. Second it is proved that a reliable scheme for encoding and decoding exists which comes arbitrarily close to *attaining* the upper bound found in 1.

This maximum rate at which information can be reliably sent through the channel is known as the *channel capacity*.

In this Dissertation we consider only the first of these tasks, the placing of upper bounds on the rate at which quantum information can be reliably sent through a noisy quantum channel, with high dynamic fidelity the criterion for successful transmission of quantum information. That is, we place bounds on the entropy of the source states, ρ_s, that can be reliably sent through such a channel.

The results we will prove are analogous to the weak converse of the classical noisy coding theorem, but cannot be considered true converses, since we do not prove that our bounds can be achieved by a coding scheme. Thus our results cannot be considered to be a channel capacity theorem, although if attainability of the upper bounds we prove could be shown, then a true channel capacity theorem would result. I do consider the bounds to be likely candidates for the quantum channel capacity.

10.4.1 Mathematical formulation of noisy channel coding

Up to this point the procedure for doing noisy channel coding has been discussed in broad outline, but we have not made all of our definitions mathematically precise. This subsection gives a mathematically precise formulation for the most important concepts appearing in our work on noisy channel coding.

Define a *quantum source*, $\Sigma \equiv (H_s, \Upsilon)$ to consist of a Hilbert space H_s and a sequence $\Upsilon = [\rho_s^1, \rho_s^2, ..., \rho_s^n, ...]$ where ρ_s^1 is a density operator on H_s, ρ_s^2 a density operator on $H_s \otimes H_s$, and ρ_s^n a density operator on $H_s^{\otimes n}$, etc... Using, for example, "tr_{34}" to denote the partial trace over the third and fourth copies of H_s, we require as part of our definition of a quantum source that for all j and all $n > j$,

$$\mathrm{tr}_{j+1,...,n}(\rho_s^n) = \rho_s^j, \tag{10.49}$$

that is, that density operators in the sequence be consistent with each other in the sense that earlier ones be derivable from later ones by an appropriate partial trace. The n-th density operator is meant to represent the state of n emissions from the source, normally thought of as taking n units of time. (We could have used a single density operator on a countably infinite tensor product of spaces H_s, but we wish to avoid the technical issues associated with such products.) We will define the *entropy* of a general source Σ as

$$S(\Sigma) \equiv \lim_{n\to\infty} \frac{S(\rho_s^n)}{n}, \tag{10.50}$$

when this limit exists.

A special case of this general definition of quantum source is the i.i.d. source $(H_s, [\rho_s, \rho_s \otimes \rho_s, ..., \rho_s^{\otimes n}, ...])$, for some fixed ρ_s. Such a source corresponds to the classical notion of an *independent, identically distributed* classical source, thus the term i.i.d. The entropy of this source is simply $S(\rho_s)$.

A discrete memoryless channel, $(H_c, \mathcal{N})$ consists of a finite-dimensional Hilbert space, H_c, and a trace-preserving quantum operation $\mathcal{N}$. The *nth extension* of that channel is given by the pair $(H_s^{\otimes n}, \mathcal{N}^{\otimes n})$. The memoryless nature of the channel is reflected in the fact that the operation performed on the n copies of the channel system is a tensor product of independent single-system operations.

Define an *n-code* from H_s into H_c to consist of a trace-preserving quantum operation, $\mathcal{C}$, from $H_s^{\otimes n}$ to $H_c^{\otimes n}$, and a trace-preserving quantum operation $\mathcal{D}$ from $H_c^{\otimes n}$ to $H_s^{\otimes n}$. We will refer to $\mathcal{C}$ as the *encoding* and $\mathcal{D}$ as the *decoding*.

The *total coding operation* $\mathcal{T}$ is given by

$$\mathcal{T} \equiv \mathcal{D} \circ \mathcal{N}^{\otimes n} \circ \mathcal{C}. \tag{10.51}$$

The measure of success we will use for the total procedure is the *total dynamic fidelity*,

$$F(\rho_s^n, \mathcal{T}). \tag{10.52}$$

In practice we will frequently abuse notation, usually by omitting explicit mention of the Hilbert spaces H_s and H_c. Note also that in principle the channel could have different input and output Hilbert spaces. To ease notational clutter we will not consider that case here, but all the results we prove go through without change.

Given a source state ρ_s and a channel $\mathcal{N}$, the goal of noisy channel coding is to find an encoding $\mathcal{C}$ and a decoding $\mathcal{D}$ such that $F(\rho_s, \mathcal{T})$ is close to one; that is, ρ_s and its entanglement is transmitted almost perfectly. In general this is not possible to do. However, Shannon showed in the classical context that by considering blocks of output from the source, and performing block encoding and decoding it is possible to considerably expand the class of source states ρ_s for which this is possible. The quantum mechanical version of this procedure is to find a sequence of n-codes, $(\mathcal{C}_n, \mathcal{D}_n)$ such that as $n \to \infty$, the measure of success $F(\rho_s^n, \mathcal{T}_n)$ approaches one, where $\mathcal{T}_n = \mathcal{D}_n \circ \mathcal{N}^{\otimes n} \circ \mathcal{C}_n$. (We will sometimes refer to such a sequence as a *coding scheme*.)

Suppose such a sequence of codes exists for a given source Σ. In this case the channel is said to transmit ρ_s reliably. We also say that the channel can transmit reliably at a *rate* $R = S(\Sigma)$. (Note that this definition does not require that the channel be able to transmit reliably *any* source with entropy less than or equal to R; that is a different potential definition of what it means for a channel to transmit reliably at rate R; in the contexts considered in this Chapter, it has been shown elsewhere [9] that the two to turn out to be equivalent, that is if a channel can transmit some source with entropy R, it can transmit any source with that entropy.)

A noisy channel coding theorem would enable one to determine, for any source and channel, whether or not the source can be transmitted reliably on that channel at a given rate. Classically, this is determined by comparing the Shannon entropy of the source to the capacity of the channel. If the entropy of the input distribution is greater than the capacity, the source cannot be transmitted reliably. If the entropy is less than the capacity, it can. The conjunction of these two statements is the noisy channel coding theorem. (The case of H precisely equal to C requires separate consideration; sometimes reliable transmission is achievable, and sometimes not.) We expect that in quantum mechanics, the entropy $S(\Sigma)$ of the source will play a role analogous to the Shannon entropy, and the coherent information will play a role analogous to the mutual information. A channel will be able to transmit reliably any source with von Neumann entropy less than the capacity; furthermore, *no* source with entropy greater than the capacity will be reliably transmissible. The first part of this would constitute a quantum noisy channel coding theorem; the second, a "weak converse" of the theorem. A "strong converse" would require not just that no source with entropy greater than the capacity can be reliably transmitted, that is transmitted with asymptotic fidelity approaching unity, but would require that all such sources have asymptotic fidelity of transmission approaching zero.

10.5 Upper bounds on the channel capacity

In this section we investigate a variety of upper bounds on the capacity of a noisy quantum channel.

10.5.1 Unitary encodings

This subsection will be concerned with the case where the encoding, $\mathcal{C}$, is unitary.

For this subsection only we define

$$C_n \equiv \max_{\rho} I(\rho, \mathcal{N}^{\otimes n}), \tag{10.53}$$

where the maximization is over all inputs ρ to n copies of the channel. The bound on the channel capacity proved in this section is defined by

$$C \equiv \lim_{n\to\infty} \frac{C_n}{n}. \tag{10.54}$$

It is not immediately obvious that this limit exists. To see that it does, notice that $C_n \leq n \log d$ and $C_m + C_n \leq C_{m+n}$ and apply the following lemma. Notice that $C = C(\mathcal{N})$ is a function of the noisy channel only.

Lemma 5 *Suppose $c_1, c_2, \ldots$ is a nonnegative sequence such that $c_n \leq kn$ for some $k \geq 0$, and*

$$c_m + c_n \leq c_{m+n}, \tag{10.55}$$

for all m and n. Then

$$\lim_{n\to\infty} \frac{c_n}{n} \tag{10.56}$$

exists and is finite.

Proof
Define

$$c \equiv \limsup_n \frac{c_n}{n}. \tag{10.57}$$

This always exists and is finite, since $c_n \leq kn$ for some $k \geq 0$. Fix $\epsilon > 0$ and choose n sufficiently large that

$$\frac{c_n}{n} > c - \epsilon. \tag{10.58}$$

Suppose m is any integer strictly greater than $\max(n, n/\epsilon)$. Then by (10.55),

$$\frac{c_m}{m} \geq \frac{c_n}{n}\frac{n}{m}\left(1 + \frac{c_{m-n}}{c_n}\right). \tag{10.59}$$

Using the fact that $lc_n \leq c_{ln}$, (an immediate consequence of (10.55)) with $l = \lfloor \frac{m}{n} \rfloor - 1$ gives

$$\frac{c_{m-n}}{c_n} \geq \lfloor \frac{m}{n} \rfloor - 1 \tag{10.60}$$

$$\geq \frac{m}{n} - 2, \tag{10.61}$$

where $\lfloor x \rfloor$ is the integer immediately below x. Plugging the last inequality into (10.59) gives

$$\frac{c_m}{m} \geq \frac{c_n}{n}\left(1 - \frac{n}{m}\right). \tag{10.62}$$

But $-n/m > -\epsilon$ and $c_n/n \geq c - \epsilon$, so

$$\frac{c_m}{m} \geq (c - \epsilon)(1 - \epsilon). \tag{10.63}$$

This equation holds for all sufficiently large m, and thus

$$\liminf_n \frac{c_n}{n} \geq (c - \epsilon)(1 - \epsilon). \tag{10.64}$$

But ϵ was an arbitrary number greater than 0, so letting $\epsilon \to 0$ we see that

$$\liminf_n \frac{c_n}{n} \geq c = \limsup_n \frac{c_n}{n}. \tag{10.65}$$

It follows that $\lim_n c_n/n$ exists, as claimed.
QED

The following theorem places a limit on the entropy of a source which can be sent through a quantum channel.

Theorem 27 ***(Upper bound on the capacity with unitary encodings)***
Suppose we consider a source $\Sigma = (H_s, [..\rho_s^n...])$ and a sequence of unitary encodings $\mathcal{U}_n$ for the source. Suppose further that there exists a sequence of decodings, $\mathcal{D}_n$ such that

$$\lim_{n\to\infty} F(\rho_s^n, \mathcal{D}_n \circ \mathcal{N}^{\otimes n} \circ \mathcal{U}_n) = 1. \tag{10.66}$$

Then

$$\limsup_{n\to\infty} \frac{S(\rho_s^n)}{n} \leq C. \tag{10.67}$$

Proof

What this theorem tells us is that we cannot reliably transmit more than C qubits of information per use of the channel. When the source entropy exists, it tells us we cannot transmit sources with entropy greater than C; when the entropy of the source is not defined, it still rules out transmission of sources for which the limsup in the expression (which is always defined) is too large. For unitary $\mathcal{U}_n$ we have

$$I(\rho_s, \mathcal{N}^{\otimes n} \circ \mathcal{U}_n) = I(\mathcal{U}_n(\rho_s), \mathcal{N}^{\otimes n}), \tag{10.68}$$

and thus

$$I(\rho_s, \mathcal{N}^{\otimes n} \circ \mathcal{U}_n) \leq C_n. \tag{10.69}$$

By (10.21) with $\mathcal{E} \equiv \mathcal{N}^{\otimes n} \circ \mathcal{U}_n$, and the fact that $I(\rho^{\otimes n}, \mathcal{N}^{\otimes n}) \leq \max_{\rho^n} I(\rho^n, \mathcal{N}^{\otimes n}) \equiv C_n$ it now follows that

$$\begin{aligned} \frac{S(\rho_s^n)}{n} \quad &\leq \quad \frac{C_n}{n} + \frac{2}{n} + \\ &\qquad 4(1 - F(\rho_s^n, \mathcal{D}^n \circ \mathcal{N}^n \circ \mathcal{U}^n)) \log d. \end{aligned} \tag{10.70}$$

(Note that d here is the dimension of a single copy of the source Hilbert space, so that we have inserted d^n for the overall dimension d of (10.21)). Taking lim sups on both sides of the equation completes the proof of the theorem.

QED

It is extremely useful to study this result at length, since the basic techniques employed to prove the bound are the same as those that appear in a more elaborate guise later in the Chapter. It is particularly instructive to see how this result differs from the classical result. In particular, what features of quantum mechanics necessitate a change in the proof methods used to obtain the classical bound?

Suppose the quantum analogue of the classical subadditivity of mutual information were true, namely

$$I(\rho^n, \mathcal{N}^{\otimes n}) \leq \sum_{i=1}^{n} I(\rho_i^n, \mathcal{N}), \tag{10.71}$$

where ρ^n is any density operator that can be used as input to n copies of the channel, and ρ_i^n is the density operator obtained by tracing out all but channel number i. Then it would follow easily from the definition that $C_n = C_1$ for all n, and thus

$$C = C_1 = \max_{\rho} I(\rho, \mathcal{N}). \tag{10.72}$$

This expression is exactly analogous to the classical expression for channel capacity as a maximum over input distributions of the mutual information between channel input and output. If this were truly a bound on the quantum channel capacity then it would allow easy numerical evaluations of bounds on the channel capacity, as the maximization involved is easy to do numerically, and the coherent information is not difficult to evaluate.

Unfortunately, it is not possible to assume that the quantum mechanical coherent information is subadditive, as shown by example (10.42), and thus in general it is possible that

$$C > C_1. \tag{10.73}$$

In fact, the results of Shor and Smolin [166] demonstrate the existence of channels for which the above strict inequality holds. In order to evaluate the bound C which we have derived it is thus necessary to take the limit in (10.54). To numerically evaluate this limit directly is certainly not a trivial task, in general. The result we have presented, that (10.54) is an upper bound on channel capacity is an important theoretical result, that may aid in the development of effective numerical procedures for obtaining general bounds. But it does not yet constitute an effective procedure.

10.5.2 General encodings

We will now consider the case where something more general than a unitary encoding is allowed. In principle, it is always possible to perform a non-unitary encoding, $\mathcal{C}$, by introducing an extra ancilla system, performing a joint unitary on the source plus ancilla, and then discarding the ancilla.

We define

$$C^n \equiv \max_{\rho,\mathcal{C}} I(\rho, \mathcal{N}^n \circ \mathcal{C}), \tag{10.74}$$

where the maximization is over all inputs ρ to the encoding operation, $\mathcal{C}$, which in turn maps to n copies of the channel,

$$\mathcal{N}^n \equiv \mathcal{N} \otimes \ldots \otimes \mathcal{N}; \quad n\text{times}, \tag{10.75}$$

The bound on the channel capacity proved in this section is defined by

$$C(\mathcal{N}) \equiv \lim_{n\to\infty} \frac{C_n}{n}. \tag{10.76}$$

Once again, to prove that this limit exists one applies the lemma proved on page 200.

To prove that this quantity is a bound on the channel capacity, one applies almost exactly the same reasoning as in the preceding subsection. The result is:

Theorem 28 ***(General bound on the channel capacity)***

Suppose we consider a source $\Sigma = (H_s, [..\rho_s^n...])$ and a sequence of unitary encodings $\mathcal{U}_n$ for the source. Suppose further that there exists a sequence of decodings, $\mathcal{D}_n$ such that

$$\lim_{n\to\infty} F(\rho_s^n, \mathcal{D}^n \circ \mathcal{N}^{\otimes n} \circ \mathcal{C}^n) = 1. \tag{10.77}$$

Then

$$\limsup_{n\to\infty} \frac{S(\rho_s^n)}{n} \leq C. \tag{10.78}$$

Proof

Again, this result places an upper bound on the rate at which information can be reliably transmitted through a noisy quantum channel. The proof is very similar to the earlier proof of a bound for unitary encodings. One simply applies (10.21) with $\mathcal{E} = \mathcal{N}^{\otimes n} \circ \mathcal{C}^n$ and $\mathcal{D} = \mathcal{D}^n$, again invoking the fact that $I(\rho^{\otimes n}, \mathcal{N}^{\otimes n}) \leq \max_{\rho^n} I(\rho^n, \mathcal{N}^{\otimes n})$, and chooses $\mathcal{C}_n$ to be the coding that maximizes this expression, to give:

$$\frac{S(\rho_s^n)}{n} \leq \frac{C_n}{n} + \frac{2}{n} + 4(1 - F(\rho_s^n, \mathcal{D}^n \circ \mathcal{N}^{\otimes n} \circ \mathcal{C}^n)) \log d. \tag{10.79}$$

Taking lim sups on both sides of the equation completes the proof.

QED

It is instructive to see why the proof fails when the maximization is done over channel input states alone, rather than over all source states and encoding schemes. The basic idea is that there may exist source states, ρ_s, and encoding schemes $\mathcal{C}$, for which

$$I(\rho, \mathcal{N} \circ \mathcal{C}) > I(\mathcal{C}(\rho), \mathcal{N}). \tag{10.80}$$

It is clear that the existence of such a scheme would cause the line of proof suggested above to fail. Moreover, as we saw in subsection 10.3.3, it is possible for exactly this situation to occur, due to quantum entanglement.

Having proved that $C(\mathcal{N})$ is an upper bound on the channel capacity, let us now investigate some of the properties of this bound. First of all we will examine the range over which C can vary. Note that

$$0 \leq C_n \leq n \log d, \tag{10.81}$$

since if ρ is pure then $I(\rho, \mathcal{N}^n \circ \mathcal{C}) = 0$ for any encoding $\mathcal{C}$, and for all ρ and $\mathcal{C}$, $I(\rho, \mathcal{N}^n \circ \mathcal{C}) \leq \log d^n = n \log d$, since the channel output has d^n dimensions. It follows that

$$0 \leq C(\mathcal{N}) \leq \log d. \tag{10.82}$$

This parallels the classical result, which states that the channel capacity varies between 0 and $\log s$, where s is the number of channel symbols. The upper bound on the classical capacity is attained if and only if the classical channel is noiseless.

In the case when $\mathcal{N}$ takes a constant value,

$$\mathcal{N}(\rho) = \sigma, \tag{10.83}$$

for all channel inputs ρ, it is not difficult to verify that $C(\mathcal{N}) = 0$. This is consistent with the obvious fact that the capacity for coherent quantum information of such a channel is zero.

When is the upper bound, $C(\mathcal{N}) = \log d$ attained? Suppose the channel is unitary, $\mathcal{N}(\rho) = U\rho U^\dagger$. Encoding the source $\rho_s = I/d \otimes \ldots \otimes I/d$ using the identity encoding, we see that $I(\rho_s, \mathcal{N}^n \circ \mathcal{C}) = \log d$, and thus $C_n \geq n \log d$, and thus $C(\mathcal{N}) \geq \log d$. But the reverse inequality also holds as remarked earlier, and thus

$$C(\mathcal{N}) = \log d, \tag{10.84}$$

if $\mathcal{N}$ is a unitary channel.

It is also of interest to consider what happens when channels $\mathcal{N}_1$ and $\mathcal{N}_2$ are composed, forming a joint channel, $\mathcal{N} = \mathcal{N}_2 \circ \mathcal{N}_1$. From the data processing inequality it follows that

$$C(\mathcal{N}_1) \geq C(\mathcal{N}). \tag{10.85}$$

It is clear by repeated application of the data-processing inequality that this result also holds if we compose more than two channels together, and even holds if we allow intermediate decoding and re-encoding stages. Classically, channel capacities also behave in this way: the capacity of a channel made by composing two (or more) channels together is no greater than the capacity of the first part of the channel alone.

Although (10.30) might seem to suggest otherwise, in fact

$$C(\mathcal{N}_2) \geq C(\mathcal{N}). \tag{10.86}$$

For let us suppose that $\mathcal{C}$ is the encoding which achieves the channel capacity $C(\mathcal{N})$, so that the total operation is $\mathcal{D} \circ \mathcal{N} \circ \mathcal{C} \equiv \mathcal{D} \circ \mathcal{N}_2 \circ \mathcal{N}_1 \circ \mathcal{C}$. As our encoding for the channel $\mathcal{N}_2$, we may use $\mathcal{N}_1 \circ \mathcal{C}$ and decode with $\mathcal{D}$, hence achieving precisely the same total operation.

10.5.3 Other encoding protocols

So far we have considered two allowed classes of encodings: encodings where a general unitary operation can be performed on a block of qubits, and encodings where a general trace-preserving quantum operation can be performed on a block of qubits. If large-scale quantum computation ever becomes feasible it may be realistic to consider encoding protocols of this sort. However, for present-day applications of quantum communication such as quantum cryptography and teleportation only a much more restricted class of encodings is possible. In this section we will describe several plausible classes.

We will begin by considering a toy example which is meant to illustrate the basic techniques which will be used later. It is the class involving local unitary operations only. We will refer to this class as U-L. It consists of the set of operations $\mathcal{C}$ which can be written in the form

$$\mathcal{C}(\rho) = (U_1 \otimes \ldots \otimes U_n)\rho(U_1^\dagger \otimes \ldots U_n^\dagger), \tag{10.87}$$

where $U_1, \ldots, U_n$ are local unitary operations on systems 1 through n. Another possibility is the class L of encodings involving local operations only, i.e. operations of the form:

$$\sum_{i_1,\ldots i_N} (A_{i_1} \otimes B_{i_2} \otimes \cdots \otimes Z_{i_N})\rho \\ (A_{i_1}^\dagger \otimes B_{1_2}^\dagger \otimes \cdots \otimes Z_{i_N}^\dagger). \tag{10.88}$$

In other words, the overall operation has a tensor product form $\mathcal{A}_1 \otimes \mathcal{A}_2 \otimes \cdots \otimes \mathcal{A}_m$.

A more realistic class is 1-L – encoding by local operations with one way classical communication. The idea is that the encoder is allowed to do encoding by performing arbitrary quantum operations on individual members (typically, a single qubit) of the strings of qubits emitted by a source. This is not unrealistic with present day technology for manipulating single qubits. Such operations could include arbitrary unitary rotations, and also generalized measurements. After the qubit is encoded, the results of any measurements done during the encoding may be used to assist in the encoding of later qubits. This is what we mean by one way communication - the results of the measurement can only be used to assist in the encoding of later qubits, not earlier qubits.

Another possible class is 2-L - encoding by local operations with two-way classical communication. These may arise in a situation where there are many identical channels operating side by side in space. Once again it is assumed that the encoder can perform arbitrary local operations, only this time two way classical communication is allowed when performing the encoding.

For any class of encodings Λ arguments analogous to those used above for general and for unitary block coding, ensure that the capacity

$$C_\Lambda(\mathcal{N}) \equiv \lim_{n\to\infty} \frac{C_\Lambda^n}{n}, \tag{10.89}$$

where

$$C_\Lambda^n \equiv \max_{\rho,\mathcal{C}\in\Lambda} I(\rho, \mathcal{N}^n \circ \mathcal{C}), \tag{10.90}$$

is an upper bound to the rate at which information can be reliably transmitted using encodings in Λ. Thus we have expressions for C_U, C_L, C_{1-L}, and C_{2-L}, which provide upper bounds on the rate of quantum information transmission for these types of encodings.

An interesting and important question is whether there are closed-form characterizations of the sets of quantum operations corresponding to particular types of encodings schemes such as 1-L

and 2-L. For example, in the cases of U-L and L there are explicit forms (10.87,10.88) for the classes of encodings allowed. For 1-L we believe the operations take the form:

$$\sum_{i_1,\ldots i_N} (A_{i_1} \quad \otimes \quad B_{i_1,i_2} \otimes \cdots \otimes Z_{i_1,i_2,\ldots i_N})\rho \\ (A_{i_1}^\dagger \otimes B_{i_1,i_2}^\dagger \otimes \cdots \otimes Z_{i_1,i_2,\ldots i_N}^\dagger). \tag{10.91}$$

It would be valuable to limit the range of the indices in this expression. This is likely to be related to the number of rounds of classical communication which are involved in an operation. Since communication is one-way, it is likely this is bounded. It would also be useful to find a similar expression for 2-L encodings. One possibility is:

$$\sum_i (A_i \otimes B_i \otimes \cdots \otimes Z_i)\rho(A_i^\dagger \otimes B_i^\dagger \otimes \cdots \otimes Z_i^\dagger). \tag{10.92}$$

However, although all 2-L operations involving a finite number of rounds of communication can certainly be put in this form, I do not presently see whether all operations expressible in this form should be realizable with local operations and two-way classical communication.

Such closed-form expressions would aid in numerical maximizations like that performed in calculating of bounds on the channel capacity. In order to perform such maximizations it would be necessary that the closed form expressions be *bounded* in size (hence the interest in limiting the range of indices above).

The classes we have described in this subsection are certainly not the only realistic classes of encodings. Many more classes may be considered, and in specific applications this may well be of great interest. What we have done is illustrated a general technique for obtaining *bounds* on the channel capacity for different classes of encodings. A major difference between classical information theory and quantum information theory is the greater interest in the quantum case in studying different classes of encodings. Classically it is, in principle, easy to perform an arbitrary encoding and decoding operation using a look-up table. However, quantum mechanically this is far from being the case, so there is correspondingly more interest in studying the channel capacities that may result from considering different classes of encodings.

Here we have not addressed the attainability of the bounds we have described. To qualify as true quantum capacities one must exhibit explicit coding and decoding schemes which allow the bounds described in this section to be achieved. The development of general proofs showing that this can be done or counterexamples showing that it cannot is a major remaining goal of quantum information theory.

10.6 Discussion

What then can be said about the status of the quantum coherent noisy channel coding theorem in the light of comments made in the preceding sections? While we have established upper bounds, we have not proved achievability. Lloyd [118] also proposed the maximum of the coherent information as the channel capacity, although initially without considering the difficulties engendered by the failure of subadditivity. He argued that this capacity was an achievable upper bound on transmission rate. The methods by which we derived our upper bound are quite different from those employed by Lloyd; I hope comparison of the two approaches will prove illuminating. The fidelity criterion he used, average pure-state fidelity for the uniform ensemble over the typical subspace, is different from the criterion used here, and although I think it is likely that they lead to the same capacity asymptotically, I

am not aware of results that imply this. In Lloyd's work, although the encoding scheme is not explicitly written out, it appears to be restricted to projection onto the typical subspace followed by a unitary. However, one can still make progress towards a proof that the novel expression, (10.76), which we have shown bounds the channel capacity, is in fact the true capacity of a noisy quantum channel for sending coherent quantum information. If we accept Lloyd's claim that his expression for the channel capacity is correct for the case when only restricted encodings are allowed, then it is possible to use the following four-stage construction to show that (10.76) is a correct expression for the capacity for transmission through a noisy quantum channel; i.e. that in addition to being an upper bound as shown in section 10.5, it is also achievable.

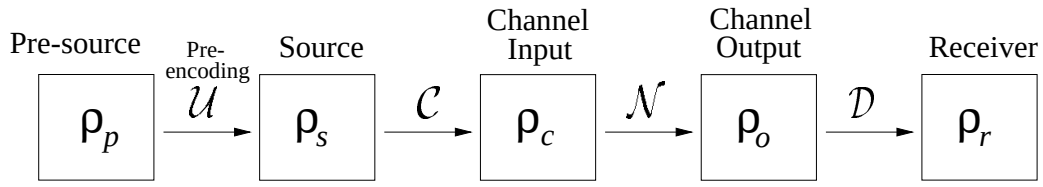

Figure 10.4: Noisy quantum channel with an extra stage, a restricted *pre-encoding*, $\mathcal{P}$.

For a fixed block size, n, one finds an encoding, $\mathcal{C}_n$, for which the maximum in

$$C_n \equiv \max_{\mathcal{C}_n, \rho_s} I(\rho_s, \mathcal{C}_n), \tag{10.93}$$

is achieved. One then regards the composition $\mathcal{N}^{\otimes n} \circ \mathcal{C}_n$ as a single noisy quantum channel, and applies the achievability result on restricted encodings to the joint channel $\mathcal{N}^{\otimes n} \circ \mathcal{C}_n$ to achieve an even longer mn block coding scheme with high dynamic fidelity. This gives a joint coding scheme $\mathcal{P}_{mn} \circ \mathcal{C}_n^{\otimes m}$ which for sufficiently large blocks m and n can come arbitrarily close to achieving the channel capacity (10.76).

An important open question is whether (10.76) is equal to (10.54). It is clear that the former expression is at least as large as the latter. Work in progress [9] shows that this is, in fact, the case.

Thus, I think it likely that the expression (10.54) will turn out to be the maximum achievable rate of reliable transmission through a quantum channel. But this is still not quite as satisfactory as the classical expression for the capacity, because of the difficulty of evaluating the limit involved. At a minimum, we would like to know enough about the rate of convergence of C_n to its limit to be able to accurately estimate the error in a numerical calculation of capacity, giving an effective procedure for calculating the capacity to any desired degree of accuracy.

10.7 Channels with a classical observer

In this section we consider a more general version of the quantum noisy channel coding problem than has been considered in any previous work. Suppose that in addition to a noisy interaction with the environment there is also a classical observer who is able to perform a measurement. This measurement may be on the channel or the environment of the channel, or possibly on both.

The result of the measurement is then sent to the decoder, who may use the result to assist in the decoding. We will assume to begin that this transmission of classical information is done noiselessly, although it is also interesting to consider what happens when the classical transmission also involves noise. It can be shown [102] that the state received by the decoder is again related to

the state ρ used as input to the channel by a quantum operation $\mathcal{N}_m$, where m is the measurement result recorded by the classical observer,

$$\rho \to \frac{\mathcal{N}_m(\rho)}{\mathrm{tr}(\mathcal{N}_m(\rho))}. \tag{10.94}$$

The basic situation is illustrated in figure 10.5. The idea is that by giving the decoder access to classical information about the environment responsible for noise in the channel it may be possible to improve the capacity of that channel, by allowing the decoder to choose different decodings $\mathcal{D}_m$ depending on the measurement result m.

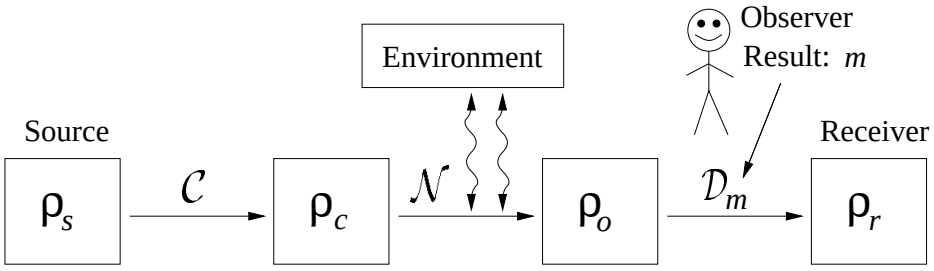

Figure 10.5: Noisy quantum channel with a classical observer.

A simple example which illustrates that this can be the case will now be given. Suppose have a two-level system in a state ρ and an initially uncorrelated four-level environment initially in the maximally mixed state $I/4$, so the total state of the joint system is

$$\rho \otimes \frac{I}{4}. \tag{10.95}$$

Suppose we fix an orthonormal basis $|1\rangle, |2\rangle, |3\rangle, |4\rangle$ for the environment. We assume that a unitary interaction between the system and environment takes place, given by the unitary operator

$$U = I \otimes |1\rangle\langle 1| + X \otimes |2\rangle\langle 2| + Y \otimes |3\rangle\langle 3| + Z \otimes |4\rangle\langle 4|. \tag{10.96}$$

The output of the channel is thus

$$\rho \to \mathcal{N}(\rho) \equiv \mathrm{tr}_E(U\rho \otimes \frac{I}{4}U^\dagger). \tag{10.97}$$

The quantum operation $\mathcal{N}$ can be given two particularly useful forms,

$$\mathcal{N}(\rho) = \frac{1}{4}\left(I\rho I + X\rho X + Y\rho Y + Z\rho Z\right) \tag{10.98}$$

$$= \frac{I}{2}. \tag{10.99}$$

It is not difficult to show from the second form that

$$C(\mathcal{N}) = 0. \tag{10.100}$$

Suppose now that an observer is introduced, who is allowed to perform a measurement on the environment. We will suppose this measurement is a Von Neumann measurement in the $|1\rangle, |2\rangle, |3\rangle, |4\rangle$ basis, and yields a corresponding measurement result, $m = 1, 2, 3, 4$. Then the quantum operations corresponding to these four measurement outcomes are

$$\mathcal{N}_1(\rho) = \frac{1}{4}\rho \tag{10.101}$$

$$\mathcal{N}_2(\rho) = \frac{1}{4}X\rho X \tag{10.102}$$

$$\mathcal{N}_3(\rho) = \frac{1}{4}Y\rho Y \tag{10.103}$$

$$\mathcal{N}_4(\rho) = \frac{1}{4}Z\rho Z. \tag{10.104}$$

Each of these is unitary, up to a constant multiplying factor, so the corresponding channel capacities are

$$C_m = 1. \tag{10.105}$$

Thus $0 = C < C_m = 1$ for each result m. Clearly this is consistent with the fact that $\mathcal{N} = \sum_m \mathcal{N}_m$ (cf. (10.17)).

This result is particularly clear in the context of teleportation. In section 3.3 we showed that the problem of teleportation can be understood precisely as the problem of a quantum noisy channel with an auxiliary classical channel. In the original single qubit teleportation scheme described in section 2.3 [18] there are four quantum operations relating the state Alice wishes to teleport, to the state Bob receives, corresponding to each of the four measurement results. In that scheme it happens that those four operations are the $\mathcal{N}_m$ we have described above. Furthermore in the absence of the classical channel, that is, when Alice does not send the result of her measurement to Bob, the channel is described by the single operation $\mathcal{N}$. Clearly, in order that causality be preserved we expect that $C = 0$. On the other hand, in order that teleportation be able to occur we should expect that $C_m = 1$, as was shown above. Teleportation understood in this way as a noisy channel with a classical side channel offers a particularly elegant way of seeing that the transmission of quantum information may sometimes be greatly improved by making use of classical information.

The remainder of this section is organized into two subsections. Subsection 10.7.1 proves bounds on the capacity of an observed channel. These results require nontrivial extension of the techniques developed earlier for proving bounds on the capacity of an unobserved channel. Subsection (10.7.2) relates work done on the observed channel to the work done earlier in the Chapter on the unobserved channel.

10.7.1 Upper bounds on channel capacity

As for the unobserved channel we will now prove several results bounding the channel capacity of an observed channel. We begin with the key lemma that will be used to prove bounds.

The following lemma generalizes the entropy-fidelity lemma on page 194 for quantum operations, which was the foundation of our earlier proofs of upper bounds on the quantum channel capacity:

Lemma 6 ***(generalized entropy-fidelity lemma for operations)***

Suppose $\mathcal{E}_m$ are a set of quantum operations such that $\sum_m \mathcal{E}_m$ is a complete quantum operation. Suppose further that $\mathcal{D}_m$ is a complete quantum operation for each m. Then

$$\begin{aligned} S(\rho) &\leq \sum_m tr(\mathcal{E}_m(\rho)) I(\rho, \mathcal{E}_m) + 2 + \\ & \quad 4(1 - F(\rho, \mathcal{T})) \log d, \end{aligned} \tag{10.106}$$

where

$$\mathcal{T} \equiv \sum_m \mathcal{D}_m \circ \mathcal{E}_m. \tag{10.107}$$

Proof

By the second step of the data processing inequality, (9.24), $I(\rho, \mathcal{E}_m) \geq I(\rho, \mathcal{D}_m \circ \mathcal{E}_m)$ for each m, and noting also that by the completeness of $\mathcal{D}_m$, $\mathrm{tr}(\mathcal{E}_m(\rho)) = \mathrm{tr}((\mathcal{D}_m \circ \mathcal{E}_m)(\rho))$, we obtain

$$\begin{aligned} S(\rho) &\leq S(\rho) + \sum_m \left[\mathrm{tr}(\mathcal{E}_m(\rho)) I(\rho, \mathcal{E}_m) - \right. \\ & \quad \left. \mathrm{tr}((\mathcal{D}_m \circ \mathcal{E}_m)(\rho)) I(\rho, \mathcal{D}_m \circ \mathcal{E}_m) \right]. \end{aligned} \tag{10.108}$$

Applying now the convexity theorem for coherent information,

$$-\sum_m \mathrm{tr}((\mathcal{D}_m \circ \mathcal{E}_m)(\rho)) I(\rho, \mathcal{D}_m \circ \mathcal{E}_m) \leq -I(\rho, \mathcal{T}). \tag{10.109}$$

we obtain

$$S(\rho) \leq \sum_m \mathrm{tr}(\mathcal{E}_m(\rho)) I(\rho, \mathcal{E}_m) + S(\rho) - I(\rho, \mathcal{T}). \tag{10.110}$$

But $\mathcal{T} = \sum_m \mathcal{D}_m \circ \mathcal{E}_m$ is trace-preserving since $\mathcal{D}_m$ is trace-preserving and $\sum_m \mathcal{E}_m$ is trace-preserving, and thus by (9.81),

$$S(\rho) - I(\rho, \mathcal{T}) = S(\rho) - S(\mathcal{T}(\rho)) + S(\rho, \mathcal{T}) \tag{10.111}$$

$$\leq 2S(\rho, \mathcal{T}). \tag{10.112}$$

Finally, an application of the quantum Fano inequality (9.11) along with the observations that the entropy function h appearing in that inequality is bounded above by one, and $\log(d^2 - 1) \leq 2 \log d$ gives

$$\begin{aligned} S(\rho) &\leq \sum_m \mathrm{tr}((\mathcal{D}_m \circ \mathcal{E}_m)(\rho)) I(\rho, \mathcal{D}_m \circ \mathcal{E}_m) + 2 + \\ & \quad 4(1 - F(\rho, \mathcal{T})) \log d, \end{aligned} \tag{10.113}$$

as we set out to prove.

QED

If we define

$$\begin{aligned} C(\{\mathcal{N}_m\}) = \limsup_{n \to \infty} \max_{\mathcal{C}^n, \rho} \\ \sum_{m_1, \ldots m_n} \mathrm{tr}((\mathcal{C}^n \circ \mathcal{N}_{m_1} \otimes \cdots \otimes \mathcal{N}_{m_n})(\rho)) \frac{I(\rho, \mathcal{N}_{m_1} \otimes \cdots \otimes \mathcal{N}_{m_n} \circ \mathcal{C}^n)}{n}, \end{aligned} \tag{10.114}$$

we may use (10.106) to easily prove that $C(\{\mathcal{N}_m\})$ is an upper bound on the rate of reliable transmission through an observed channel, in precisely the same way we earlier used (10.21) to prove bounds for unobserved channels.

We may derive the same bound in another fashion if we associate observed channels with complete quantum operations – unobserved channels – in the following fashion suggested by examples in [21]. To an observed channel $\{\mathcal{N}_m\}$ we associate a single complete operation $\mathcal{M}$ from $\mathcal{H}$ to the larger Hilbert space $\mathcal{H} \otimes \mathcal{R}$. The operation is specified by:

$$\mathcal{M}(\rho) = \sum_m \mathcal{N}_m(\rho) \otimes |m\rangle\langle m|. \tag{10.115}$$

This map is an "all-quantum" version of the observed channel. The classical information about which m occurred appears in the "register" Hilbert space $\mathcal{R}$ encoded in orthogonal states. Since our upper bound to the capacity of an unobserved channel applies also to channels with output Hilbert spaces of different dimensionality than the input space, they apply to this map as well. It is easily verified that the coherent information for the map $\mathcal{M}$ acting on ρ is the same as the average coherent information for the observed channel $\mathcal{N}_m$ acting on ρ, which appears in (10.106) and in the quantity (10.114). To show this, define $p_m = \mathrm{tr}(\mathcal{N}_m(Q))$. Then $Q' = \mathcal{M}(Q)$ is given by (10.115), so that

$$S(Q') = H(p_m) + \sum_m p_m S\left(\frac{\mathcal{N}_m(Q)}{p_m}\right) \tag{10.116}$$

by the grouping property (9.14) of Shannon entropy, which applies since the density matrices $\mathcal{N}_m(Q) \otimes |m\rangle\langle m|$ are mutually orthogonal. Similarly,

$$R'Q' = (\mathcal{I} \otimes \sum_m \mathcal{N}_m^*)(RQ), \tag{10.117}$$

where by definition $\mathcal{N}_m^*(\rho) = \mathcal{N}_m(\rho) \otimes |m\rangle\langle m|$. By linearity this is equal to $\sum_m (\mathcal{I} \otimes \mathcal{N}_m)(RQ) \otimes |m\rangle\langle m|$. Applying the orthogonality and grouping argument again, and noting that $\mathrm{tr}((\mathcal{I}\otimes\mathcal{N}_m)(RQ)) = \mathrm{tr}(\mathcal{N}_m(Q)) = p_m$, we get that

$$S(R'Q') = H(p_m) + \sum_m p_m S\left(\frac{(\mathcal{I} \otimes \mathcal{N}_m)(RQ)}{p_m}\right). \tag{10.118}$$

Hence the coherent information for $\mathcal{M}$ becomes

$$\sum_m p_m \left[S\left(\frac{\mathcal{N}_m(Q)}{p_m}\right) - S\left(\frac{(\mathcal{I} \otimes \mathcal{N}_m)(RQ)}{p_m}\right)\right], \tag{10.119}$$

which is precisely the average coherent information for $\{\mathcal{N}_m\}$. So an application of the bound (10.76) on the rate of transmission through the unobserved channel $\mathcal{M}$ yields the bound (10.114), if one accepts the intuitively obvious claim that $\mathcal{M}$ and $\{\mathcal{N}_m\}$ are equivalent with respect to transmission of quantum information.

Bennett *et al* [21] derive capacities for three simple channels which may be viewed as taking the form (10.115). The *quantum erasure channel* takes the input state to a fixed state orthogonal to the input state with probability ϵ; otherwise, it transmits the state undisturbed. An equivalent observed channel would with probability ϵ replace the input state with a standard pure state $|0\rangle\langle 0|$ within the input subspace, and also provide classical information as to whether this replacement has occurred or not. The *phase erasure channel* randomizes the phase of a qubit (or, in our context of

multidimensional input space, diagonalizes the density operator in a fixed basis) with probability delta, and otherwise transmits the state undisturbed; it also supplies classical information as to which of these alternatives occurred. The *mixed erasure/phase-erasure channel* may either erase or phase-erase, with exclusive probabilities ϵ and δ. Bennett *et al* note that the capacity $1-2\epsilon$ of the erasure channel is in fact the one-shot maximal coherent information. We have verified that the capacities they derive for the phase-erasure channel $(1-\delta)$ and the mixed erasure/phase-erasure channel $(1-2\epsilon-\delta)$ are the same as the one-shot maximal average coherent information for the corresponding observed channels, lending some additional support to the view that the bounds we have derived here are in fact the capacities.

10.7.2 Relationship to unobserved channel

Suppose we have an observed channel which is described by operations $\{\mathcal{N}_m\}$. There are various natural physical ways these operations can be associated with a channel described by the operation

$$\mathcal{N} = \sum_m \mathcal{N}_m. \tag{10.120}$$

One physically natural way this association may be made is the following. Suppose a system is sent through a noisy quantum channel. During the time and possibly after the system has traversed the channel, various measurements may be performed, possibly on the system, and possibly on the environment giving rise to the noise. We will label the collective results of these measurements by a single index, m. As discussed earlier, with each m is associated a quantum operation, $\mathcal{N}_m$, which describes the state change undergone by a system passing through the channel, given that result m occurs. (If the measurement involves the system, and not just the environment, there is no guarantee that $\mathcal{N} \equiv \sum_m \mathcal{N}_m = \mathcal{N}_0$, where $\mathcal{N}_0$ is the noise due to the channel without measurement.)

There is a particularly important special case of the above scenario. Suppose the system is sent through a channel, and interacts with an environment. The action of this channel is described by the complete quantum operation $\mathcal{N}_0$. After the environment has interacted with the system, measurements are performed on the *environment alone*. Averaging over all possible measurement outcomes, this does not disturb the state of the system, i.e. $\mathcal{N} \equiv \sum_m \mathcal{N}_m = \mathcal{N}_0$.

We will now show that observing the environment of the channel never *decreases* the bound we have obtained on the channel capacity. This is certainly a property which we would expect the channel capacity to have: observing the environment and then sending the result of the observation on to be used in decoding should not decrease the channel capacity, since the decoder can always simply ignore the result of the observation.

Recall the expressions for the bound on the capacity of the unobserved channel,

$$C(\mathcal{N}) = \limsup_{n\to\infty} \max_{\mathcal{C}^n,\rho} \frac{I(\rho, \mathcal{N}^{\otimes n} \circ \mathcal{C}^n)}{n}, \tag{10.121}$$

and the observed channel,

$$C(\{\mathcal{N}_m\}) = \limsup_{n\to\infty} \max_{\mathcal{C}^n,\rho} \sum_{m_1,\ldots m_n} \mathrm{tr}((\mathcal{C}^n \circ \mathcal{N}_{m_1} \otimes \cdots \otimes \mathcal{N}_{m_n})(\rho)) \frac{I(\rho, \mathcal{N}_{m_1} \otimes \cdots \otimes \mathcal{N}_{m_n} \circ \mathcal{C}^n)}{n}. \tag{10.122}$$

But the convexity theorem for coherent information implies that

$$\sum_{m_1,\ldots,m_n} \mathrm{tr}((\mathcal{C}^n \circ \mathcal{N}_{m_1} \otimes \cdots \otimes \mathcal{N}_{m_n})(\rho)) \frac{I(\rho, \mathcal{N}_{m_1} \otimes \cdots \otimes \mathcal{N}_{m_n} \circ \mathcal{C}^n)}{n}$$

$$\leq \frac{I(\rho, \mathcal{N} \circ \mathcal{C}^n)}{n}, \tag{10.123}$$

and thus

$$C(\mathcal{N}) \leq C(\{\mathcal{N}_m\}). \tag{10.124}$$

To see that this inequality may sometimes be strict, return to the example considered earlier in the context of teleportation. In that case it is not difficult to verify that

$$0 = C(\mathcal{N}) < C(\{\mathcal{N}_m\}) = 1. \tag{10.125}$$

What these results show is that our bounds on the channel capacity are never made any worse by observing the environment, but sometimes they can be made considerably better. This is a property that we certainly expect the quantum channel capacity to have, and we take as an encouraging sign that the bounds we have proved in this Chapter are in fact achievable, that is, the true capacities.

10.7.3 Discussion

All the questions asked about the bounds on channel capacity for an unobserved channel can be asked again for the observed channel: questions about achievability of bounds, the differences in power achievable by different classes of encodings and decodings, and so on. We will not address those problems here, beyond noting that they are important problems which need to be addressed by future research.

Many new twists on the problem of the quantum noisy channel arise when an observer of the environment is allowed. For example, one might consider the situation where the classical channel connecting the observer to the decoder is noisy. What then are the resources required to transmit coherent quantum information?

It might also be interesting to prove results relating the classical and quantum resources that are required to perform a certain task. For example, in teleportation it can be shown that one requires not only the quantum channel, but also two bits of classical information, in order to transmit coherent quantum information with perfect reliability [18].

10.8 Conclusion

In this Chapter we have shown that different information transmission problems may result in different channel capacities for the same noisy quantum channel. We have developed some general techniques for proving upper bounds on the amount of information that may be transmitted reliably through a noisy quantum channel.

Perhaps the most interesting thing about the quantum noisy channel problem is to discover what is new and essentially *quantum* about the problem. The following list summarizes what I believe are the essentially new features:

1. The insight that there are many essentially different information transmission problems in quantum mechanics, all of them of interest depending on the application. These span a spectrum between two extremes:

   - The transmission of a discrete set of mutually orthogonal quantum states through the channel. Such problems are problems of transmitting classical information through a noisy quantum channel.

- The transmission of entire subspaces of quantum states through the channel, keeping entanglement intact. This is likely to be of interest in applications such as quantum computation, cryptography and teleportation where superpositions of quantum states are crucial. Such problems are problems of transmitting coherent quantum information through a noisy quantum channel.

Both these cases are important for specific applications. For each case, there is great interest in considering different classes of allowed encodings and decodings. For example, it may be that encoding and decoding can only be done using local operations and one-way classical communication. This may give rise to a different channel capacity than occurs if we allow non-local encoding and decoding. Thus there are different noisy channel problems depending on what class of encodings and decodings is allowed.

2. The use of quantum entanglement to construct examples where the quantum analogue of the classical equation $H(X : Z) \leq H(Y : Z)$ for a Markov process $X \rightarrow Y \rightarrow Z$, fails to hold (compare equation (10.30)).

3. The use of quantum entanglement to construct examples where the subadditivity property of mutual information,

$$H(X_1, \ldots, X_n : Y_1, \ldots, Y_n) \leq \sum_i H(X_i : Y_i), \tag{10.126}$$

fails to hold (compare equation (10.42)).

There are many more interesting open problems associated with the noisy channel problem than have been addressed here. The following is a sample of those problems which I believe to be particularly important:

1. The development of good numerical algorithms for determining the different channel capacities. If the expressions for channel capacities involve limits like those in the upper bounds in this Chapter, it will also be important either to evaluate those limits analytically, or to know the rate of convergence to those limits to aid in evaluating them numerically.

2. Estimation of channel capacities for realistic channels. This work could certainly be done theoretically and perhaps also experimentally, using the technique of quantum process tomography discussed in section 3.4. An interesting problem is to analyze how stable the determination of channel capacities is with respect to experimental error.

3. As suggested in subsection 10.5.3 it would be interesting to see what channel capacities are attainable for different classes of allowable encodings, for example, encodings where the encoder is only allowed to do local operations and one-way classical communication, or encodings where the encoder is allowed to do local operations and two-way classical communication. We have seen how to prove bounds on the channel capacity in these cases; whether these bounds are attainable is unknown.

4. The development of rigorous general techniques for proving attainability of channel capacities, which may be applied to different classes of allowed encodings and decodings.

5. Finding the capacity of a noisy quantum channel for classical information; considerable progress on this problem has already been made [85, 155], however much remains to be done. A related problem arises in the context of *superdense coding*, where one half of an EPR pair can be used

to send two bits of classical information. It would be interesting to know to what extent this performance is degraded if the pair of qubits shard between sender and receiver is not an EPR pair, but rather the sharing is done using a noisy quantum channel, leading to a decrease in the number of classical bits that can be sent. Given a noisy quantum channel, what is the maximum amount of classical information that can be sent in this way?

6. All work done thus far has been for discrete channels, that is, channels with finite dimensional state spaces. It is an important and non-trivial problem to extend these results to channels with infinite dimensional state spaces.

There are many other ways the classical results on noisy channels have been extended - considering channels with *feedback*, developing *rate-distortion* theory and so on. Each of these could give rise to highly interesting work on noisy quantum channels. It is also to be expected that interesting new questions will arise as experimental efforts in the field of quantum information develop further. My own chief interest to us is to develop a still clearer understanding of the essential differences between the quantum noisy channel and the classical noisy channel problem, and to provide an effective procedure for evaluating and achieving the quantum channel capacity.

Summary of Chapter 10: The quantum channel capacity

- **The quantum channel capacity:** The maximum rate at which quantum information can be sent through a noisy quantum channel.
- **Coherent information:**

$$I(\rho, \mathcal{E}) \equiv S(\rho') - S(\rho, \mathcal{E}).$$

  Behaves in many ways as a quantum analogue to the classical mutual information.
- **Failure of the data pipelining inequality for the coherent information:**

$$I(\rho, \mathcal{E}_2 \circ \mathcal{E}_1) \not\leq I(\mathcal{E}_1(\rho), \mathcal{E}_2).$$

  The extent to which this inequality is violated is a lower bound on the entanglement between E_1'' and E_2''.
- **Failure of subadditivity for the coherent information:**

$$I(\rho_{12}, \mathcal{E} \otimes \mathcal{E}) \not\leq I(\rho_1, \mathcal{E}) + I(\rho_2, \mathcal{E}).$$

- **General bound on quantum channel capacity:**

$$C(\mathcal{N}) \leq \lim_{n \rightarrow \infty} \max_{\rho^n, \mathcal{C}^n} \frac{I(\rho^n, \mathcal{N}^{\otimes n} \circ \mathcal{C}^n}{n}.$$

  It can be shown [9] that the maximization over encodings $\mathcal{C}^n$ is unnecessary, and can be removed.
- **The observed quantum channel:** By performing measurements upon the environment of a quantum system we may be able to use the result of the measurement to increase the quantum channel capacity.

Part III

Conclusion

Chapter 11

Conclusion

This concluding Chapter briefly surveys some future directions in quantum information theory. We begin by summarizing the results of the Dissertation, with an emphasis on novel results, and new problems suggested by this research. We conclude by broadening our net to look at the wider field of quantum information, suggesting some possibly profitable directions for future research.

11.1 Summary of the Dissertation

We begin by reviewing what has been achieved in the Dissertation, and open problems arising directly as a result of this research. The major achievement of the Dissertation is the discovery of numerous bounds on our ability to do quantum information processing, and the development of new techniques for proving such bounds.

Part I of the Dissertation reviewed some of the tools necessary to make progress in quantum information theory, especially, the quantum operations formalism, the properties of entropy in quantum mechanics, and distance measures for quantum information. In many ways it was gaining a good understanding of these tools that I expect to be the most useful aspect of doing the research that led to this Dissertation over the long term, although the research results in Part I are of a fairly diffuse nature; the primary purpose of Part I is pedagogical.

Part II of the Dissertation consisted largely of original research on problems in quantum information theory, focusing on the proof of bounds to what tasks are possible in quantum information theory, using the tools introduced in Part I of the Dissertation.

In Chapter 6 we investigated quantum communication complexity, the study of the communication requirements involved in distributed quantum computation. Holevo's theorem was used to prove a new capacity theorem which encapsulates the limits to communication of classical information between two parties when a two-way noiseless quantum channel is available for use between the parties. A generalization of this result to the case of channels with noise would be of great interest. The capacity theorem was then applied to prove that the availability of a noiseless quantum channel does not assist in the calculation of the inner product of two bit strings, when one of those bit strings belongs to one party, and the other bit string belongs to a second party. This is a significant result, as it tells us that there are problems in communication complexity for which quantum mechanics provides no advantage over the classical result. The Chapter also contained the first results in coherent quantum communication complexity, which deals with the communication requirements incurred when computing a *unitary operator* in a distributed fashion. We were able to show that the computation of the quantum Fourier transform over $2n$ qubits, n of them belonging to one party,

and n to a second party, requires at least n qubits of communication between those parties. A general lower bound on the coherent communication complexity was proved, which we may expect to be of great assistance in future investigations of the coherent communication complexity. Finally, the beginnings of a framework which unifies previous work on quantum communication complexity was sketched. Most importantly, it includes as special cases both the coherent communication complexity, and the communication complexity of a classical function, using quantum resources.

In Chapter 7 we studied the compression of information from a quantum source. A new proof was given of the quantum data compression theorem, which gives an *operational interpretation* of the von Neumann entropy $S(\rho)$ as the minimal number of qubits with which it is possible to reliably store a quantum source described by the density operator ρ. The new techniques introduced during the proof were then applied to the problem of *universal quantum data compression*, a hoped-for technique which provides a means for compressing quantum information even in the absence of knowledge about a quantum source's characteristics. We constructed a method for performing a potentially useful form of universal quantum data compression on large class of quantum sources, although it remains to find an efficient quantum algorithm for implementing this procedure for universal quantum data compression.

Chapter 8 focused on the central problem of providing good quantitative measures of entanglement. Entanglement appears to be a central resource in most quantum information processing tasks known to date. This Chapter focused on one particular measure of entanglement, the *entanglement of formation*, a measure of how many Bell states it takes to create a particular entangled state. The most important result of the Chapter was a relationship between the entanglement of formation and the quantum conditional entropy, $E(A:B) \geq -S(A|B)$. In particular, this shows that the puzzling phenomena of negative quantum conditional entropies are always associated with the presence of entanglement in a quantum system.

Chapter 9 introduced the basic notions of quantum error correction. After reviewing the basic ideas of quantum error correction using the Shor nine qubit code, we developed information-theoretic necessary and sufficient conditions for quantum error correction, using a quantum analogue of the classical data processing inequality. Next, we analyzed the thermodynamic cost of quantum error correction, showing that quantum error correction schemes function as a kind of Maxwell's demon, in which information is extracted from a system in order to lower its entropy. We were able to show that quantum error correction can be performed in a thermodynamically efficient manner.

Chapter 10 studied the problem of the quantum channel capacity. The channel capacity measures how much quantum information can be sent through a noisy quantum channel. We developed a bound on this quantity based upon the coherent information, and explained some of the outstanding problems related to the channel capacity. The problem of the *observed quantum channel* was introduced, and we explained how it can be reformulated in purely quantum terms. Upper bounds on the channel capacity of an observed quantum channel were proved, again using the coherent information.

What are the most important outstanding problems arising directly from the dissertation? Perhaps the most immediately fruitful areas for future research are quantum communication complexity and the understanding of entanglement.

With regard to quantum communication complexity, the general lower bound technique for quantum communication complexity which I proved in Chapter 6 can doubtless be generalized, and applied to many interesting problems. It would, for example, be useful to understand the communication costs in performing quantum Fourier transforms over groups more general than the integers modulo 2^n, as we considered, or to consider the communication costs incurred when doing quantum error correction. More generally, little is presently known about the relationship between quantum and classical communication complexity. Developing these connections more deeply is an

obvious area for further work.

The study of entanglement is a second area that I believe will yield rich results over the next few years. Among many possible avenues of research, I am especially interested in pursuing the behaviour of entanglement in systems in thermodynamic equilibrium, and trying to understand how the entanglement behaves near some of the phase transitions which may occur in such systems.

The other topics addressed by the Dissertation – quantum data compression, quantum error correction, and the channel capacity – also suggest many interesting problems. Developing a more complete understanding of universal quantum data compression is a worthwhile goal, and may be of practical importance in the future. I am also actively investigating the problem of data compression of correlated quantum sources in an attempt to find a quantum analogue of the Slepian-Wolf theorem of classical information theory [54]; these results are incomplete, and were not included in the Dissertation. With regard to error correction and the channel capacity, perhaps the most outstanding problem is to better understand the quantum channel capacity, and to develop a general procedure for evaluating it, analogous to Shannon's noisy channel coding theorem. This is a fascinating, albeit apparently quite a difficult problem, whose solution I expect will give us great new insight into quantum information.

Summarizing, in this Dissertation I have discovered many new limits to the ability to perform information processing within quantum mechanics. Many of these limitations are concerned with multiple parties, where some bound is placed on their communication requirements. These result in practical limits on the ability of two parties to compress quantum data, to communicate classical data using quantum resources, to compute the inner product function and the quantum Fourier transform, to perform quantum error correction, and to send quantum information through a noisy quantum channel. More general theoretical results have been proved which give general though not always practically applicable bounds on the ability to perform distributed quantum computations, and on the ability to send quantum information through a noisy quantum channel. Moreover, new tools and techniques have been developed while solving these problems that will be of great use in further investigations of quantum information theory. Perhaps most importantly, though, many interesting new problems have been raised. We now turn to look more broadly at the problems facing quantum information theory at the present time.

11.2 Open problems in quantum information

In the last section we reviewed the achievements of this Dissertation, and some of the open problems arising directly from this research. In this section we discuss in a broader setting some of the challenges facing quantum information theory. Several simply stated problems may be identified as especially important:

1. Develop computationally interesting new applications of quantum information.

2. What are the ultimate achievable limits to quantum information processing? Several subproblems may be identified:

   - What class of problems may be solved efficiently on a quantum computer? How does this class compare to the class of problems efficiently soluble on a classical computer?
   - What resources are required to do distributed quantum computation?
   - In a multi-party situation where not all parties trust one another, what resources are required to do information processing tasks with a reasonable level of security?

3. What technologies are needed to implement quantum information processing? Does it have economically practical applications?

4. Can quantum information shed new light on the problems of fundamental physics?

5. Can other physical theories be used to do information processing tasks beyond the quantum computational model?

One of the fun and exciting things about quantum information is that we are still at the point where simple, fundamental questions like these can be asked, without the answers being known. Precious little is known about the computational power of quantum mechanics. Remedying this situation offers many interesting challenges.

A discussion of directions to take in solving the listed problems would be enormous. Instead, I will focus on three directions which I believe are especially promising as first steps if we are to solve these problems.

The first of these directions is inspired directly by the results in this Dissertation. In subsection 11.2.1 I will sketch out a formalism which can be used to unify several of the disparate approaches to quantum information theory which have been developed. In particular, results such as the Holevo theorem which relate to sending classical resources through quantum channels, appear to have little relationship to questions of manipulating quantum information in quantum channels. I sketch out the beginnings of a formalism which can unite these approaches to quantum information theory.

The second direction which we look at is the so-called "decoherence program", which aims to explain how classical physics arises as the limit to quantum physics. Throughout this Dissertation, we have assumed that there are two fundamental units of information, one classical (the "bit"), one quantum (the "qubit"). It appears as though Nature does not respect such a duality at the fundamental level. Rather, classical information arises as the limit of quantum information under certain special circumstances. Understanding how this occurs in more detail is the goal of the decoherence program. In subsection 11.2.2 I ask whether quantum information theory can information this program.

The third direction which we will examine is whether there are interesting cross-disciplinary advances to be made between quantum information theory and statistical physics. This is the most speculative and the most sketchy of all the proposals made here. Nevertheless, I feel it is a direction well worth exploring.

11.2.1 A unifying picture for quantum information

This Dissertation has explored many parts of quantum information theory. Despite considerable effort on my part, I did not find it possible to present the results of all chapters within a single, unified picture. Compare, for example, Chapter 7, on quantum data compression, with Section 6.1, on the Holevo bound. Although similar tools are used in each instance, it is not immediately apparent that the approaches to both problems fit within a single, unified approach to quantum information theory. Nowhere is this lack of a unified approach more apparent than in the existence of two different pictures, one to handle the problem of sending classical information, using quantum resources, the other to handle the problem of sending quantum information using quantum resources.

This is a problem which exists throughout the published literature on quantum information. Recently, an interesting paper by Schumacher and Westmoreland [156] has appeared which demonstrates a link between the two approaches to quantum information theory. Broadly speaking, Schumacher and Westmoreland draw our attention to a connection between a quantity based on

the Holevo χ which measures how much *classical information* can be sent, with *complete privacy*, through a quantum channel, and the *coherent information*, which we have seen is related to how much *quantum information* can be sent through a channel.

This work has stimulated me to think about whether a unified picture for the description of quantum information might be found. I believe I have found the beginnings of such a picture, which I will sketch in this section. Nevertheless, considerable work remains to be done before the new picture can be considered complete. A partial statement of the goals of such a picture is that it ought to be able to achieve *all* of the following in an integrated manner:

1. Describe an ensemble of (potentially mixed) states being produced by a classical source.
2. Describe a source which represents an entanglement with a reference system.
3. Describe (potentially multi-part) dynamical processes.
4. Describe the classical results of (possibly non-ideal) measurements.
5. Describe a closed universe in which the state of the total system is pure at all times.
6. Describe how classical information theory and the classical world arise as a limit of quantum information theory and the quantum world.

This seems like a rather lengthy list of features to require in a single picture! The picture I describe will seem somewhat complicated. Nevertheless, it offers a remarkably simple way for rederiving results such as the Schumacher and Westmoreland result.

Suppose a classical source is producing quantum states ρ_x according to some probability distribution p_x. Define $\rho \equiv \sum_x p_x \rho_x$. The following construction enables us to describe the classical source, together with an entangled source producing the state ρ, all within the one formalism.

Let $Q \equiv Q_1$ be the quantum system under consideration. Let Q_2 be a copy of system Q, and let the states $|Q_1 Q_2\rangle_x$ be purifications of the states ρ_x. Under some circumstances it may be advantageous to require additional properties of the purifications $|Q_1 Q_2\rangle_x$[1], but we will not need such additional properties here.

We also introduce two systems, P_1 and P_2, each of which has an orthonormal basis $|x\rangle$ of states in one-to-one correspondence with the outputs which may be produced by the classical source. Defining $P \equiv P_1$, we will refer to the system P as the *preparation system*, since it will be used to encode information about which of the state ρ_x has been prepared.

Notationally, we will write states of the system $P_1 P_2 Q_1 Q_2$ in the order $P_1 P_2 Q_1 Q_2$, unless otherwise noted. The state of the system representing both the classical source $\{p_x, \rho_x\}$ and the entangled source ρ is given by

$$|P_1 P_2 Q_1 Q_2\rangle = \sum_x \sqrt{p_x} |x\rangle |x\rangle |Q_1 Q_2\rangle_x. \tag{11.1}$$

Note that requirement number 5 is met – this state is pure.

Note that

$$PQ = \sum_x p_x |x\rangle\langle x| \otimes \rho_x. \tag{11.2}$$

[1]For example, we might enlarge Q_2, and required that the states $|Q_1 Q_2\rangle_x$ be mutually orthogonal. Or we might require that the reduced states on system Q_2, Q_{2x}, are equal to the original states ρ_x of Q_1. Other possibilities may also be useful.

Intuitively, the system PQ therefore describes a classical system P which is in one of the orthonormal states $|x\rangle$, with respective probabilities p_x, and which prepares a corresponding state of system Q, ρ_x. Compare this with the similar construction in the proof of Holevo's bound, section 6.1. That is, requirement 1 is met.

To see that requirement 2 is met, define the *reference system* $R \equiv P_1 P_2 Q_2$. Notice that RQ is a purification of the system Q, which starts in the state $Q = \sum_x p_x \rho_x$.

Requirements 3 and 4 require that we introduce additional systems. Suppose $\mathcal{E}$ is a complete quantum operation that may occur on the system. This quantum operation may arise, potentially, as the result of a measurement described by quantum operations $\mathcal{E}_m$, $\mathcal{E} = \sum_m \mathcal{E}_m$. Suppose $\mathcal{E}_m$ has an operator-sum representation generated by operators E_{mi}. Introduce a system M with an orthonormal basis $|m\rangle$, a system I with an orthonormal basis $|i\rangle$, and a system E with an orthonormal basis $|m,i\rangle$. Let $|0_M\rangle, |0_I\rangle, |0_E\rangle$ be standard pure states of the respective systems M, I and E. Define a unitary operator U on $QMIE$ which has the action

$$U|\psi\rangle|0\rangle|0\rangle|0\rangle \equiv \sum_{mi} E_{mi}|\psi\rangle|m\rangle|i\rangle|m,i\rangle. \tag{11.3}$$

Intuitively, the system M plays the part of a "measuring apparatus", which records the result of the measurement. The system I represents information which is lost when the measurement is performed. Finally, the system E can be thought of as an "environment", which decoheres the measuring apparatus, in the sense of Zurek [206].

In the language used to prove the Holevo bound in section 6.1, M plays the same role here as it does there – it stores the result of the measurement. The joint system IE plays the same role as the system E did in section 6.1. Finally, the combined system MIE plays the role of a model environment for the operation $\mathcal{E} = \sum_m \mathcal{E}_m$, just as was used throughout the Dissertation.

Thus, this formalism encompasses all the constructions contained in this Dissertation, from constructions in which there is a classical source of information, as in the Holevo bound, through to results such as the data processing inequality, which deal with the transmission of entanglement through a quantum channel. Indeed, part of the great attraction of this formalism is the number of results which it gives you *automatically*. In section 9.6 we were able to obtain numerous entropy inequalities, simply by applying the subadditivity and strong subadditivity inequality in a mindless fashion. In a similar way, one can obtain the Holevo bound, the data processing inequality, the result of Schumacher and Westmoreland linking the Holevo bound and coherent information [156], and many other results, all free-of-charge, as automatic consequences of the formalism and a few powerful tools such as strong subadditivity. For this reason, I believe that this formalism will be a powerful tool to aid in answering new questions about the connection between classical and quantum information.

11.2.2 Classical physics and the decoherence program

Since the foundation of quantum mechanics there has been intense interest in understanding how the classical world we see in our everyday life arises out of the underlying quantum reality. This effort was particularly intense in the early years, with researchers such as Bohr [27] stressing general principles which linked quantum and classical physics, and researchers such as Mott [128] who did detailed investigations of specific phenomena in order to explain how the classical world arises from the quantum reality.

Since those early days there has been a continual effort to understand the connection between quantum reality and the classical world. Unfortunately, this work suffered considerably because of a lack of experimental progress at the level of single quantum systems. With a few notable exceptions,

much of the work in this area got bogged down in a morass of theory and philosophy, with only a few notable pieces of science emerging from sixty-odd years of work.

Over the past twenty years or so there has been tremendous experimental progress at the level of single quantum systems. To cite but one important example, there was the development in the mid 1980s of the *quantum jumps* technique, a technique for doing projective measurements on a single ion in an ion trap [129, 150, 24], based upon a Technical Report written by Dehmelt in the mid 1970s, but never published.

This experimental Renaissance in studying foundational issues in quantum mechanics has been matched by a similar theoretical Renaissance. A particularly broad program, sometimes known under the rubric of the "decoherence program" has been advanced by Zurek and other researchers, starting in the early 1980s [202, 203]. Reviews of this material may be found in [208, 206]. The decoherence program is an attempt to give a detailed explanation for how classical behaviour arises from quantum reality. Much of this problem is now well understood, at least in outline, but some fundamental problems remain.

How may the tools of quantum information be brought to bear upon the decoherence program? I do not know how to answers this question, but I do have a number of problems which I would like solved:

1. What does it mean to say that system A "objectively knows" something about system B? Answering this question in a quantitative fashion ought to give us a much better handle on the old problem of determining, from first principles, what physical systems function as measuring devices, in the sense of inducing a collapse of the state vector.

2. Zurek [207] has proposed what he refers to as the "predictability sieve". This is a proposal intended to solve the following fundamental problem: given a unitary interaction between a quantum system and a measuring device, there are many possible "collapse" rules for the quantum system consistent with the non-selective unitary interaction between system and measuring device. Zurek gives a prescription for determining which of the collapse rules consistent with the unitary interaction is actually taking place. This prescription uses the von Neumann entropy to determine the "most classical" set of quantum states possessed by the measuring apparatus. Unfortunately, the motivation for using the von Neumann entropy in this context has never been physically clear. It would be interesting to approach the predictability sieve from the point of view of quantum information theory, and to ask what is the appropriate quantity for measuring the "classicallity" of a set of states of the measuring device.

3. Zurek [208] has asked why it is that a composite, "system" structure exists in nature? This structure is crucial to the success of the decoherence program, but it has never been explained why it is found in nature. It is not obvious that this question is immediately related to quantum information theory, however, it is as well to keep what is perhaps the most significant unsolved problem of the decoherence program in mind while pursuing an information-theoretic approach to decoherence.

11.2.3 Quantum information and statistical physics

An observation that has recurred in my thoughts for several years now is that statistical physics and the theory of computation seem like two different approaches to a very similar problem. In both cases we are trying to determine the long-time behaviour of a dynamical system whose constituent parts behave according to simple rules.

This observation is, apparently, not without foundation in fact, for I have recently learned that there are, in fact, deep connections which can be made between computation and statistical

physics. In particular, it has been shown that certain problems in statistical physics are NP-complete – a computer science term for a class of problems that is believed to be intractable, at least within classical computational models [67]. The class of NP-complete problems is tremendously important, containing as it does many of the most important problems in computer science. All NP-complete problems are, essentially, equivalent in terms of computational difficulty.

It is interesting to ask, then, whether similar results hold relating difficult problems for quantum computers to problems in quantum statistical mechanics? How difficult is it to predict the long-time behaviour of a quantum system such as a spin glass? What are the transport properties of entanglement in such a system? Can we relate the existence of phase transitions in quantum statistical mechanics to entanglement between the constituent parts? Answering questions such as these has the potential benefit of not just illuminating aspects of quantum information theory, but also other areas of physics.

11.3 Concluding thought

Quantum information, and more generally, the physics of information, offers tremendous opportunities. We may be able to harness the laws of physics to perform fantastic computations impossible within the classical laws. It may even be that computing devices harnessing the full power of physics will be able to achieve a fuller comprehension of the world than our classically-limited minds. The physics of information stimulates us to ask new questions about the foundations of the physical world we live in. The exploration of these possibilities is a deep and beautiful problem, full of challenges to stimulate and awe the mind that contemplates them, the hand that realizes them, and any being that makes practical use of their eventual fruits.

Appendix A

Purifications and the Schmidt decomposition

Composite quantum systems are used throughout this Dissertation. In order to get a better grasp of the properties of composite systems we need tools to understand the states of composite quantum systems. Two of the most useful tools for doing this are the *Schmidt decomposition*, and *purifications*. In this Appendix we will review both these tools, and try to give a flavour of their power. The first part of the Appendix gives a review of these results in their standard form; the second part of the Appendix gives a new generalization of the Schmidt decomposition for which, unfortunately, I have not yet been able to find any interesting applications.

Theorem 29 *(***Schmidt decomposition)** *[152]*

Suppose $|AB\rangle$ is a pure state of a composite system, AB. Then there exists orthonormal states $|i_A\rangle$ for system A, and orthonormal states $|i_B\rangle$ of system B such that

$$|AB\rangle = \sum_i \lambda_i |i_A\rangle |i_B\rangle, \tag{A.1}$$

where λ_i are positive real numbers satisfying $\sum_i \lambda_i^2 = 1$.

This innocuous looking theorem is tremendously useful. As a taste of its power, consider the following consequence: let $|AB\rangle$ be a pure state of a composite system, AB. Then the eigenvalues of A and B are identical, namely λ_i^2 for both density operators. Many important properties of quantum systems are completely determined by the eigenvalues of the system. For a pure state of a composite system such properties will therefore be the same for both systems. As an example, consider that the von Neumann entropy of a state is completely determined by the eigenvalues of that state. Therefore, for a pure state $|AB\rangle$ of system AB, the von Neumann entropy of systems A and B are the same.

Proof

Let A be the state of system A when system B is traced out, $A \equiv \mathrm{tr}_B(|AB\rangle\langle AB|)$. Let

$$A = \sum_i p_i |i_A\rangle\langle i_A| \tag{A.2}$$

be an orthonormal decomposition for system A. Then there exist vectors $|\psi_i^B\rangle$ in the state space of system B such that

$$|\psi\rangle = \sum_i |i_A\rangle|\psi_B^i\rangle. \tag{A.3}$$

But we know that $A = \mathrm{tr}_B(|AB\rangle\langle AB|)$, from which we deduce $\langle\psi_i^B|\psi_j^B\rangle = \delta_{ij}p_i$. Thus, we can find orthonormal $|i_B\rangle$ such that $|\psi_i^B\rangle = \sqrt{p_i}|i_B\rangle$, and thus

$$|AB\rangle = \sum_i \sqrt{p_i}|i_A\rangle|i_B\rangle. \tag{A.4}$$

Setting $\lambda_i \equiv \sqrt{p_i}$ and noting that $\sum_i \lambda_i^2 = \sum_i p_i = 1$ completes the proof.

QED

The bases $|i_A\rangle$ and $|i_B\rangle$ are called the *Schmidt bases* for A and B, respectively, and the number of non-zero values λ_i is called the *Schmidt number* for the state $|\psi\rangle$. The Schmidt number is a tremendously important property of a composite quantum system. To get some idea of why this is the case, consider the following obvious but important property: the Schmidt number is preserved under unitary transformations on system A or system B alone. To see this, notice that if $\sum_i \lambda_i|i_A\rangle|i_B\rangle$ is the Schmidt decomposition for $|\psi\rangle$ then $\sum_i \lambda_i(U|i_A\rangle)|i_B\rangle$ is the Schmidt decomposition for $U|\psi\rangle$, where U is a unitary operator acting on system A alone. Invariance properties of this type are very useful in Chapter 7.

Another useful technique in quantum information theory is *purification*. Suppose we are given a state A of a quantum system A. It is possible to introduce another system, which will call R, and *define* a *pure state* $|AR\rangle$ for the joint system AR such that $A = \mathrm{tr}_R(|AR\rangle\langle AR|)$. That is, the pure state $|AR\rangle$ reduces to A when we look at system A alone. This is a purely mathematical procedure, known as *purification*, which allows us to associate pure states with mixed states. For this reason we call system R a *reference* system: it is a fictitious system, without a direct physical significance.

To prove that purification can be done for *any* mixed state, we will construct a system R and purification $|AR\rangle$ corresponding to the state A. Suppose A has spectral decomposition, $A = \sum_i p_i|i\rangle\langle i|$. In order to purify A we introduce a system R which has the same state space as system A, and define a pure state for the combined system

$$|AR\rangle \equiv \sum_i \sqrt{p_i}|i\rangle|i\rangle. \tag{A.5}$$

We now calculate the reduced density operator for system A associated with $|AR\rangle$:

$$\begin{aligned}
\mathrm{tr}_R(|AR\rangle\langle AR|) &= \sum_{ij} \sqrt{p_i p_j}|i\rangle\langle j|\mathrm{tr}(|i\rangle\langle j|) && \text{(A.6)}\\
&= \sum_{ij} \sqrt{p_i p_j}|i\rangle\langle j|\delta_{ij} && \text{(A.7)}\\
&= \sum_i p_i|i\rangle\langle i| && \text{(A.8)}\\
&= A. && \text{(A.9)}
\end{aligned}$$

Thus $|AR\rangle$ is a valid purification of A.

Notice the close relationship of the Schmidt decomposition to purification: the procedure used to purify a mixed state is to define a pure state whose Schmidt basis is just the basis in

which the mixed state is diagonal. A related observation is that the Schmidt decomposition may be used to obtain a classification of purifications of the state A. First of all, note that the Schmidt decomposition implies that if $|AB\rangle$ is a purification of A then B must contain at least as many dimensions as A has support dimensions. Suppose $|AB\rangle$ and $|AB'\rangle$ are two purifications of A. Then by the Schmidt decomposition

$$|AB\rangle = \sum_i \sqrt{p_i}|i_A\rangle|i_B\rangle \tag{A.10}$$

$$|AB'\rangle = \sum_i \sqrt{p_i}|i_A\rangle|i'_B\rangle, \tag{A.11}$$

where p_i are the eigenvalues of A, $|i_A\rangle$ the corresponding eigenvectors, and $|i_B\rangle$ and $|i'_B\rangle$ are each a set of orthonormal vectors in system B. Since $|i_B\rangle$ and $|i'_B\rangle$ are both orthonormal sets, it follows that there exists a unitary operator U_B on system B such that $U_B|i_B\rangle = |i'_B\rangle$, and therefore

$$|AB\rangle = U_B|AB'\rangle. \tag{A.12}$$

Conversely, if $|AB\rangle$ is a purification of A and U_B is a unitary operator on B then $|AB'\rangle$ defined by $|AB'\rangle \equiv U_B|AB\rangle$ is easily verified to be a purification of A. Thus $|AB\rangle$ and $|AB'\rangle$ are purifications of A if and only if there exists a unitary operator on B relating the two states.

To conclude the Appendix, we present a new generalization of the Schmidt decomposition to mixed states of a two-part composite system. I have not found any applications of this result, which is why this is an Appendix, and not a Chapter, however I am hopeful that in the future it may prove useful.

Suppose ρ is any state of a composite system AB. For convenience we assume that A and B have the same number of dimensions; if this is not true then it can be made true by appending extra dimensions to whichever system has fewer dimensions. Suppose $\rho = \sum_k |k\rangle\langle k|$, where $|k\rangle$ is an orthogonal set, with the eigenvalues $\langle k|k\rangle$ of ρ absorbed into the normalizations of the states $|k\rangle$. Similarly, suppose $A = \sum_i |i\rangle\langle i|$ and $B = \sum_j |j\rangle\langle j|$, where $|i\rangle$ and $|j\rangle$ are orthogonal sets, with the eigenvalues of A and B absorbed into the normalizations of $|i\rangle$ and $|j\rangle$.

Our goal here is to take A and B as given density operators for systems A and B, and to derive a set of algebraic constraints on the matrices a^k, in order that ρ be an allowed density operator for the system AB, which when system A is traced out reduces to B, and when system B is traced out, reduces to A. The result we obtain gives as a special case the Schmidt decomposition for pure states. The key to doing this is to relate the sets $|i\rangle, |j\rangle$ and $|k\rangle$. Note first that each $|k\rangle$ must be a linear combination of the $|i\rangle$s and the $|j\rangle$s,

$$|k\rangle = \sum_{ij} a^k_{ij}|i\rangle|j\rangle. \tag{A.13}$$

Tracing out system B from the expression $\rho = \sum_k |k\rangle\langle k|$ we see that

$$\sum_i |i\rangle\langle i| = \sum_k \sum_{ii'j} q_j a^k_{ij}(a^k_{i'j})^*|i\rangle\langle i'|, \tag{A.14}$$

where $q_j \equiv \langle j|j\rangle$ are the eigenvalues of B. Defining Q to be a diagonal matrix with entries q_j, we can rewrite the previous equation as

$$\sum_k a^k Q (a^k)^\dagger = I. \tag{A.15}$$

Tracing out system A from the expression $\rho = \sum_k |k\rangle\langle k|$ we deduce that

$$\sum_j |j\rangle\langle j| = \sum_k \sum_{ijj'} p_i a_{ij}^k (a_{ij'}^k)^* |j\rangle\langle j'|, \tag{A.16}$$

where $p_i \equiv \langle i|i\rangle$ are the eigenvalues of A. Defining P to be a diagonal matrix with entries p_i, we see that

$$\sum_k (a^k)^\dagger P a^k = I. \tag{A.17}$$

Furthermore, we have the orthogonality relation

$$\langle k|k'\rangle = \sum_{ij} (a_{ij}^k)^* a_{ij}^{k'} p_i q_j = \mathrm{tr}((a^k)^\dagger P a^{k'} Q). \tag{A.18}$$

The trace condition $\mathrm{tr}(\rho) = 1$ is now seen to be equivalent to $\sum_k \mathrm{tr}((a^k)^\dagger P a^k Q) = 1$. However, since $\sum_k (a^k)^\dagger P a^k = I$, the trace condition is equivalent to $\mathrm{tr}(Q) = I$. This is true by assumption, since $\mathrm{tr}(B) = 1$, so the trace condition gives nothing new. We have proved the following theorem:

Theorem 30 ***(Mixed state Schmidt decomposition)***
Let A and B be given density operators on systems A and B. Let P be a matrix whose diagonal entries are the eigenvalues of A, and Q a matrix whose diagonal entries are the eigenvalues of ρ_B. Then ρ is a density operator of system AB, consistent with A and B, if and only if ρ has orthogonal decomposition

$$\rho = \sum_k |k\rangle\langle k|, \tag{A.19}$$

where the $|k\rangle$ are defined by

$$|k\rangle = \sum_{ij} a_{ij}^k |i\rangle|j\rangle, \tag{A.20}$$

and the a^k are complex matrices satisfying the conditions

$$\sum_k a^k Q (a^k)^\dagger = I \tag{A.21}$$

$$\sum_k (a^k)^\dagger P a^k = I \tag{A.22}$$

$$tr((a^k)^\dagger P a^{k'} Q) = \langle k|k'\rangle. \tag{A.23}$$

This result provides a complete characterization of the possible states ρ of AB, in terms of the density operators A and B. Let's see what it tells us in the pure state case. In this case, there is a single non-trivial value of k. We deduce that $aQa^\dagger = I$ and $a^\dagger P a = I$. Polar decompose a as $a = uh$, for some unitary u and Hermitian h. Note that $hQh = hu^\dagger P u h = I$, from which we see that $\det h \neq 0$ and $\det P \neq 0$, and thus $P = uQu^\dagger$ and $h = Q^{-1/2}$. Since P and Q are diagonal, u must be a permutation matrix. Relabeling the basis states, we have $P = Q$ and $u = I$. We then have

$$|k\rangle = \sum_{ij} h_{ij} |i\rangle|j\rangle \tag{A.24}$$

$$= \sum_i \frac{1}{\sqrt{p_i}} |i^A\rangle|i^B\rangle. \tag{A.25}$$

The pure state Schmidt decomposition follows by renormalizing the states$|i^A\rangle$ and $|i^B\rangle$, which at present satisfy $\langle i^A|i^A\rangle = \langle i^B|i^b\rangle = p_i = q_i$.

Bibliography

[1] D. S. Abrams and S. Lloyd. Simulation of many-body Fermi systems on a quantum computer. *Phys. Rev. Lett.*, 79(13):2586–2589, 1997. quant-ph/9703054.

[2] D. Aharonov and M. Ben-Or. Fault tolerant computation with constant error. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing*, pages 176–188, 1997.

[3] A. E. Allahverdyan and D. B. Saakian. Converse coding theorems for quantums source and noisy channel. *arXive e-print quant-ph/9702034*, 1997.

[4] P. W. Anderson. More is different. *Science*, 177(4047):393–396, 1972.

[5] H. Araki and M. M. Yanase. Measurement of quantum mechanical operators. *Phys. Rev.*, 120:622, 1960.

[6] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Theory of superconductivity. *Physical Review*, 108(5):1175–1204, 1957.

[7] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Physical Review A*, 52:3457–3467, 1995. quant-ph/9503016.

[8] H. Barnum, C. M. Caves, C. A. Fuchs, R. Jozsa, and B. Schumacher. Noncommuting mixed states cannot be broadcast. *Physical Review Letters*, 76(15):2828–2821, 1996. quant-ph/9511010.

[9] H. Barnum, E. Knill, and M. A. Nielsen. On quantum fidelities and channel capacities. 1998. quant-ph/9809010.

[10] H. Barnum, M. A. Nielsen, and B. W. Schumacher. Information transmission through a noisy quantum channel. *Phys. Rev. A*, 57:4153, 1998.

[11] D. Beckman, A. N. Chari, S. Devabhaktuni, and J. Preskill. Efficient networks for quantum factoring. *Physical Review A*, 54(2):1034, 1996. arXive E-print quant-ph/9602016.

[12] J. S. Bell. *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Mechanics*. Cambridge University Press, Cambridge, 1989.

[13] P. Benioff. The computer as a physical systems: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines. *J. Stat. Phys.*, 22(5):563–591, 1980.

[14] C. H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6):525–32, 1973.

[15] C. H. Bennett. Demons, engines and the second law. *Scientific American*, 295(5):108, 1987.

[16] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, New York, 1984. IEEE. Bangalore, India, December 1984.

[17] C. H. Bennett, G. Brassard, S. Breidbart, and S. Wiesner. Quantum cryptography, or unforgeable subway tokens. In D. Chaum, R. L. Rivest, and A. T. Sherman, editors, *Advances in Cryptology: Proceedings of Crypto 82*, pages 267–275, New York, 1982. Plenum Press.

[18] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. Wootters. Teleporting an unknown quantum state via dual classical and EPR channels. *Phys. Rev. Lett.*, 70:1895–1899, 1993.

[19] C. H. Bennett, G. Brassard, and N. D. Mermin. Quantum cryptography without Bell's theorem. *Physical Review Letters*, 68(5):557–559, 1992.

[20] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters. Purification of noisy entanglement and faithful teleportation via noisy channels. *Physical Review Letters*, 76:722, 1996. quant-ph/9511027.

[21] C. H. Bennett, D. P. DiVincenzo, and J. A. Smolin. Capacities of quantum erasure channels. *Physical Review Letters*, 78(16):3217–3220, 1997. quant-ph/9701015.

[22] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters. Mixed state entanglement and quantum error correction. *Phys. Rev. A*, 54:3824, 1996. quant-ph/9604024.

[23] C. H. Bennett and S. J. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Physical Review Letters*, 69(20):2881–2884, 1992.

[24] J. C. Bergquist, R. G. Hulet, W. M. Itano, and D. J. Wineland. Observation of quantum jumps in a single atom. *Physical Review Letters*, 57(14):1699–1702, 1986.

[25] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM J. Comp.*, 26(5):1411–1473, 1997. quant-ph/9701001.

[26] R. Bhatia. *Matrix analysis*. Springer-Verlag, New York, 1997.

[27] N. Bohr. The quantum postulate and the recent development of atomic theory. *Nature*, 121:580–590, 1928. Reprinted in [187].

[28] D. Boschi, S. Branca, F. De Martini, L. Hardy, and S. Popescu. Experimental realization of teleporting an unknown pure quantum state via dual classical and Einstein-Podolski-Rosen channels. *Phys. Rev. Lett.*, 80:1121–1125, 1998. quant-ph/9710013.

[29] D. Bouwmeester, J. W. Pan, K. Mattle, M. Eibl, H. Weinfurter, and A. Zeilinger. Experimental quantum teleportation. *Nature*, 390(6660):575–579, 1997.

[30] G. Brassard. Teleportation as a quantum computation. In T. Toffoli, M. Biafore, and J. Leao, editors, *PhysComp 96*, pages 48–50, Cambridge MA, 1996. New England Complex Systsems Institute.

[31] G. Brassard, S. Braunstein, and R. Cleve. *To appear in Physica D*, 1998.

[32] G. E. Bredon. *Topology and geometry*. Springer-Verlag, New York, 1993.

[33] A. R. Calderbank, E. M. Rains, P. W. Shor, and N. J. A. Sloane. Quantum error correction and orthogonal geometry. *Phys. Rev. Lett.*, 78:405–8, 1997.

[34] C. M. Caves. Quantitative limits on the ability of a Maxwell demon to extract work from heat. *Physical Review Letters*, 64(18):2111–2114, 1990.

[35] C. M. Caves. Information, entropy, and chaos. In J. J. Halliwell, J. Perez-Mercader, and W. H. Zurek, editors, *Physical Origins of Time Asymmetry*. Cambridge University Press, Cambridge, 1994.

[36] M.-D. Choi. Completely positive linear maps on complex matrices. *Linear Algebra and Its Applications*, 10:285–290, 1975.

[37] B. Chor and O. Goldreich. Unbiased bits from weak sources of randomness and probabilistic communication complexity. *SIAM J. on Comput.*, 17(2):230–261, 1988.

[38] I. L. Chuang, N. Gershenfeld, and M. Kubinec. Experimental implementation of fast quantum searching. *Physical Review Letters*, 18(15):3408–3411, 1998.

[39] I. L. Chuang, N. Gershenfeld, M. G. Kubinec, and D. W. Leung. Bulk quantum computation with nuclear-magnetic-resonance: theory and experiment. *Proc. Roy. Soc. Lond. A*, 454(1969):447–467, 1998.

[40] I. L. Chuang and M. A. Nielsen. Prescription for experimental determination of the dynamics of a quantum black box. *J. Mod. Opt.*, 44(11-12):2455–2467, 1997. quant-ph/9610001.

[41] I. L. Chuang, L. M. K. Vandersypen, X. L. Zhou, D. W. Leung, and S. Lloyd. Experimental realization of a quantum algorithm. *Nature*, 393(6681):143–146, 1998.

[42] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Phys. Rev. Lett.*, 74:4091, 1995.

[43] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi. Quantum state transfer and entanglement distribution among distant nodes in a quantum network. *Phys. Rev. Lett.*, 78(16):3221–3224, 1997. quant-ph/9611017.

[44] R. Cleve and H. Buhrman. Substituting quantum entanglement for communication. *Physical Review A*, 56(2):1201–1204, 1997. quant-ph/9704026.

[45] R. Cleve and D. P. DiVincenzo. Schumacher's quantum data compression as a quantum computation. *Phys. Rev. A*, 54:2636, 1996. quant-ph/9603009.

[46] R. Cleve, A. Ekert, C. Macciavello, and M. Mosca. Quantum algorithms revisited. *Proc. Roy. Soc. A*, 454(1969):339–354, 1998.

[47] R. Cleve, W. van Dam, M. A. Nielsen, and A. Tapp. Quantum entanglement and the communication complexity of the inner product function. *arXive e-print quant-ph/9708019*, 1997.

[48] D. Coppersmith. An approximate Fourier transform useful in quantum factoring. *IBM Research Report RC 19642*, 1994.

[49] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. The MIT Press, Cambridge, Mass., 1990.

[50] Intel Corporation. Moore's law: Changing the PC platform for another 20 years. *Online at http://developer.intel.com/solutions/archive/issue2/focus.htm*, 1998.

[51] D. G. Cory, A. F. Fahmy, and T. F. Havel. Ensemble quantum computing by NMR spectroscopy. *Proc. Natl. Acad. Sci. USA*, 94:1634–1639, 1997.

[52] D. G. Cory, W. Mass, M. Price, E. Knill, R. Laflamme, W. H. Zurek, T. F. Havel, and S. S. Somaroo. Experimental quantum error correction. *arXive e-print quant-ph/9802018*, 1998.

[53] D. G. Cory, M. D. Price, and T. F. Havel. Nuclear magnetic resonance spectroscopy: An experimentally accessible paradigm for quantum computing. *arXive e-print quant-ph/9709001*, 1997.

[54] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley and Sons, New York, 1991.

[55] I. Csiszár and J. Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Academic Press, New York, 1981.

[56] M. D. Davis. *The Undecidable*. Raven Press, Hewlett, New York, 1965.

[57] M. D. Davis and E. J. Weyuker. *Computability, Complexity, and Languages*. Academic Press, New York, 1983.

[58] D. Deutsch. Quantum theory, the Church-Turing Principle and the universal quantum computer. *Proc. R. Soc. Lond. A*, 400:97, 1985.

[59] D. Deutsch. Notes on the quantum Fourier transform. *Unpublished*, 1994.

[60] D. Dieks. Communication by EPR devices. *Physics Letters A*, 92(6):271–272, 1982.

[61] P. A. M. Dirac. *The principles of quantum mechanics 4th ed.* Oxford University Press, Oxford, 1958.

[62] K. E. Drexler. *Nanosystems: Molecular Machinery, Manufacturig, and Computation*. John Wiley and Sons, New York, 1992.

[63] A. Ekert and R. Jozsa. Quantum computation and Shor's factoring algorithm. *Rev. Mod. Phys.*, 68:1, 1996.

[64] R. Ernst, G. Bodenhausen, and A. Wokaun. *Principles of Nuclear Magnetic Resonance in One and Two Dimension*. Oxford University Press, Oxford, 1990.

[65] R. P. Feynman. *The character of physical law*. MIT Press, Cambridge, MA, 1965.

[66] R. P. Feynman. Simulating physics with computers. *Int. J. Theor. Phys.*, 21:467, 1982.

[67] M. H. Freedman. P/np, and the quantum field computer. *Proc. Natl. Acad. Sci. USA*, 95:98–101, 1998.

[68] C. A. Fuchs. *Distinguishability and Accessible Information in Quantum Theory*. PhD thesis, The University of New Mexico, Albuquerque, NM, 1996. Also arXive e-print quant-ph/9601020.

[69] N. Gershenfeld and I. L. Chuang. Bulk spin resonance quantum computation. *Science*, 275:350, 1997.

[70] D. Gottesman. Class of quantum error-correcting codes saturating the quantum Hamming bound. *Phys. Rev. A*, 54:1862, 1996.

[71] D. Gottesman. *Stabilizer Codes and Quantum Error Correction*. PhD thesis, California Institute of Technology, Pasadena, CA, 1997.

[72] D. Gottesman. Theory of fault-tolerant quantum computation. *Physical Review A*, 57(1):127–137, 1998. quant-ph/9702029.

[73] E. S. Grace. *Biotechnology unzipped: Promises and realities*. National Academy Press, 1997.

[74] D. M. Grant and R. K. Harris, editors. *Encyclopedia of nuclear magnetic resonance*. John Wiley, New York, 1996.

[75] R. B. Griffiths and C.-S. Niu. Semiclassical Fourier transform for quantum computation. *Physical Review Letters*, 76(17):3228–3231, 1996. quant-ph/9511007.

[76] G. R. Grimmett and D. R. Stirzaker. *Probability and Random Processes*. Clarendon Press, Oxford, 1992.

[77] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *28th ACM Symposium on Theory of Computation*, page 212, New York, 1996. Association for Computing Machinery.

[78] S. W. Hawking. Breakdown of predictability in gravitational collapse. *Phys. Rev. D*, 14:2460, 1976.

[79] K.-E. Hellwig and K. Kraus. Pure operations and measurements. *Commun. Math. Phys.*, 11:214–220, 1969.

[80] K.-E. Hellwig and K. Kraus. Operations and measurements. II. *Commun. Math. Phys.*, 16:142–147, 1970.

[81] Carl W. Helstrom. *Quantum Detection and Estimation Theory*. Mathematics in Science and Engineering, vol. 123. Academic Press, New York, 1976.

[82] S. Hill and W. K. Wootters. Entanglement of a pair of quantum bits. *Phys. Rev. Lett.*, 78(26):5022–5025, 1997. quant-ph/9703041.

[83] A. S. Holevo. Statistical problems in quantum physics. In Gisiro Maruyama and Jurii V. Prokhorov, editors, *Proceedings of the Second Japan–USSR Symposium on Probability Theory*, pages 104–119, Berlin, 1973. Springer-Verlag. Lecture Notes in Mathematics, vol. 330.

[84] A. S. Holevo. *Probabilistic and Statistical Aspects of Quantum Theory*. North-Holland Series in Statistics and Probability, vol. 1. North-Holland, Amsterdam, 1982.

[85] A. S. Holevo. The capacity of quantum channel with general signal states. *arXive e-print quant-ph/9611023*, 1996.

[86] R. A. Horn and C. R. Johnson. *Topics in matrix analysis*. Cambridge University Press, Cambridge, 1991.

[87] S. F Huelga, C. Macchiavello, T. Pellizzari, A. K. Ekert, M. B. Plenio, and J. I. Cirac. On the improvement of frequency stardards with quantum entanglement. *Phys. Rev. Lett.*, 79(20):3865–3868, 1997. quant-ph/9707014.

[88] R. J. Hughes, D. M Alde, P. Dyer, G. G. Luther, G. L. Morgan, and M. Schauer. Quantum cryptography. *Contemp. Phys.*, 36(3):149–163, 1995. quant-ph/9504002.

[89] J. A. Jones and M. Mosca. Implementation of a quantum algorithm to solve Deutsch's problem on a nuclear magnetic resonance quantum computer. *arXive e-print quant-ph/9801027*, 1998.

[90] J. A. Jones, M. Mosca, and R. H. Hansen. Implementation of a quantum search algorithm on a nuclear magnetic resonance quantum computer. *Nature*, 393(6683):344, 1998. quant-ph/9805069.

[91] K. R. W. Jones. Fundamental limits upon the measurement of state vectors. *Phys. Rev. A*, 50:3682–3699, 1994.

[92] R. Jozsa. Fidelity for mixed quantum states. *J. Mod. Opt.*, 41:2315–2323, 1994.

[93] R. Jozsa. Entanglement and quantum computation. *arXive e-print quant-ph/9707034*, 1997.

[94] R. Jozsa and B. Schumacher. A new proof of the quantum noiseless coding theorem. *J. Mod. Opt.*, 41:2343–2349, 1994.

[95] A. Y. Kitaev. Quantum error correction with imperfect gates. In A. S. Holevo O. Hirota and C. M. Caves, editors, *Quantum Communication, Computing, and Measurement*, pages 181–188, New York, 1997. Plenum Press.

[96] E. Knill, I. Chuang, and R. Laflamme. Effective pure states for bulk quantum computation. *Phys. Rev. A*, 57(5):3348–3363, 1998. quant-ph/9706053.

[97] E. Knill and R. Laflamme. Concatenated quantum codes. *arXive e-print quant-ph/9608012*, 1996.

[98] E. Knill and R. Laflamme. A theory of quantum error-correcting codes. *Phys. Rev. A*, 55:900, 1997. quant-ph/9604034.

[99] E. Knill and R. Laflamme. On the power of one bit of quantum information. *arXive e-print quant-ph/9802037*, 1998.

[100] E. Knill, R. Laflamme, and W. H. Zurek. Resilient quantum computation. *Science*, 279(5349):342–345, 1998. quant-ph/9702058.

[101] E. Knill, R. Laflamme, and W. H. Zurek. Resilient quantum computation: error models and thresholds. *Proc. Roy. Soc. A*, 454(1969):365–384, 1998. quant-ph/9702058.

[102] Karl Kraus. *States, Effects, and Operations: Fundamental Notions of Quantum Theory*. Lecture Notes in Physics, Vol. 190. Springer-Verlag, Berlin, 1983.

[103] I. Kremer. Quantum communication. Master's thesis, The Hebrew University of Jerusalem, 1995.

[104] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, Cambridge, 1996.

[105] R. Laflamme, E. Knill, W. H. Zurek, P. Catasti, and S. V. S. Mariappan. NMR GHZ. *arXive e-print quant-ph/9709025*, 1997.

[106] R. Landauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.*, 5:183, 1961.

[107] R. Landauer. Computation and physics: Wheeler's meaning circuit? *Found. Phys.*, 16(6), 1986.

[108] R. Landauer. Information is physical. *Physics Today*, 44(5):22–29, 1991.

[109] J. S. Langer. Nonequilibrium physics. In V. L. Fitch, D. R. Marlow, and M. A. E. Dementi, editors, *Critical problems in physics*, pages 281–306. Princeton University Press, Princeton, New Jersey, 1997.

[110] Y. Lecerf. Machines de Turing réversibles. *Comptes Rendus*, 257:2597–2600, 1963.

[111] D. Leibfried, D. M. Meekhof, B. E. King, C. Monroe, W. M. Itano, and D. J. Wineland. Experimental-determination of the motional quantum state of a trapped atom. *Physical Review Letters*, 77(21):4281–5285, 1996.

[112] U. Leonhardt. Discrete Wigner function and quantum-state tomography. *Physical Review A*, 53(5):2998–3013, 1996.

[113] D. W. Leung, M. A. Nielsen, I. L. Chuang, and Y. Yamamoto. Approximate quantum error correction can lead to better codes. *Phys. Rev. A*, 56:2567–2573, 1997.

[114] E. H. Lieb. Convex trace functions and the Wigner-Yanase-Dyson conjecture. *Advances in Mathematics*, 11:267–288, 1973.

[115] E. H. Lieb and M. B. Ruskai. Proof of the strong subadditivity of quantum mechanical entropy. *J. Math. Phys.*, 14:1938–1941, 1973.

[116] L. Ljung. *System identification: Theory for the user*. Prentice-Hall, Upper Saddle River, 1987.

[117] S. Lloyd. Universal quantum simulators. *Science*, 273:1073, 1996.

[118] S. Lloyd. The capacity of the noisy quantum channel. *Phys. Rev. A*, 56:1613, 1997.

[119] S. Lloyd. Quantum controllers for quantum systems. *arXive e-print quant-ph/9703042*, 1997.

[120] M. Luby. *Pseudorandomness and cryptographic applications*. Princeton University Press, Princeton, New Jersey, 1996.

[121] H. Mabuchi. Dynamical identification of open quantum-systems. *Quant. Semi. Opt.*, 8(6):1103–1108, 1996.

[122] H. Mabuchi. Standard quantum limits for broadband position measurement. *arXive e-print quant-ph/9801039*, 1998.

[123] M. Marcus and H. Minc. *A Survey of Matrix Theory and Matrix Inequalities*. Dover, New York, 1992.

[124] K. Mattle, H. Weinfurter, P. G. Kwiat, and A. Zeilinger. Dense coding in experimental quantum communication. *Physical Review Letters*, 76(25):4656–4659, 1996.

[125] M. McLuhan. *Understanding media: The extension of man*. McGraw-Hill, New York, 1964.

[126] G. J. Milburn. A quantum mechanical Maxwell's demon. Unpublished, 1996.

[127] C. Monroe, D. Leibfried, B. E. King, D. M. Meekhof, W. M. Itano, and D. J. Wineland. Simplified quantum logic with trapped ions. *Physical Review A*, 55(4):R2489–R2491, 1997.

[128] N. F. Mott. The wave mechanics of α-ray tracks. *Proc. Roy. Soc. Lond. A*, 126:79–84, 1929. Reprinted in [187].

[129] W. Nagourney, J. Sandberg, and H. Dehmelt. Shelved optical electron amplifier: observation of quantum jumps. *Physical Review Letters*, 56(26):2797–2799, 1986.

[130] A. Newell and H. Simon. Computer science as empirical enquiry: Symbols and search. *Communications of the Association for Computing Machinery*, 19:113–126, 1976.

[131] M. A. Nielsen. The entanglement fidelity and quantum error correction. *arXive e-print quant-ph/9606012*, 1996.

[132] M. A. Nielsen. Computable functions, quantum measurements, and quantum dynamics. *Phys. Rev. Lett.*, 79(15):2915–2918, 1997.

[133] M. A. Nielsen and C. M. Caves. Reversible quantum operations and their application to teleportation. *Phys. Rev. A*, 55(4):2547–2556, 1997.

[134] M. A. Nielsen, C. M. Caves, B. Schumacher, and H. Barnum. Information-theoretic approach to quantum error correction and reversible measurement. *Proc. Roy. Soc. A*, 454(1969):277–304, 1998.

[135] M. A. Nielsen and I L. Chuang. Programmable quantum gate arrays. *Phys. Rev. Lett.*, 79(2):321–324, 1997.

[136] M. A. Nielsen, E. Knill, and R. Laflamme. *http://wwwcas.phys.unm.edu∼mnielsen/nmr/index.html*, 1998.

[137] M. A. Nielsen, E. Knill, and R. Laflamme. Experimental quantum teleportation by NMR. *In preparation*, 1998.

[138] M. Ohya and D. Petz. *Quantum entropy and its use*. Springer-Verlag, Berlin, 1993.

[139] C. M. Papadimitriou. *Computational complexity*. Addison-Wesley, Reading, Massachusetts, 1994.

[140] A. Peres. *Quantum Theory: Concepts and Methods*. Kluwer Academic, Dordrecht, 1993.

[141] S. Popescu and D. Rohrlich. Thermodynamics and the measure of entanglement. *Phys. Rev. A*, 56(5):R3319–R3321, 1997.

[142] J. F. Poyatos, J. I. Cirac, and P. Zoller. Complete characterization of a quantum process: the two-bit quantum gate. *Phys. Rev. Lett.*, 78(2):390–393, 1997.

[143] J. Preskill. Fault-tolerant quantum computation. *arXive e-print quant-ph/9712048*, 1997.

[144] J. Preskill. Reliable quantum computers. *Proc. Roy. Soc. A: Math., Phys. and Eng.*, 454(1969):385–410, 1998.

[145] W. H. Press, S. A. Teukolsky, B. P. Flannery, and W. T. Vetterling. *Numerical recipes in Fortran*. Cambridge University Press, Cambridge, 1989.

[146] M. G. Raymer, M. Beck, and D. F. McAlister. Complex wave-field reconstruction using phase-space tomography. *Physical Review Letters*, 72(8):1137–1140, 1994.

[147] A. Royer. Reduced dynamics with initial correlations, and time-dependent environment and Hamiltonians. *Physical Review Letters*, 77(16):3272–3275, 1996.

[148] M. B. Ruskai. Extremal properties of relative entropy in quantum statistical mechanics. *Rep. Math. Phys.*, 26(1):143–150, 1988.

[149] M. B. Ruskai. Beyond strong subadditivity: improved bounds on the contraction of generalized relative entropy. *Rev. Math. Phys.*, 6(5A):1147–1161, 1994.

[150] T. Sauter, W. Neuhauser, R. Blatt, and P. E. Toschek. Observation of quantum jumps. *Physical Review Letters*, 57(14):1696–1698, 1986.

[151] R. Schack. Algorithmic information and simplicity in statistical physics. *arXive e-print hep-th/9409022*, 1994.

[152] E. Schmidt. Zur theorie der linearen und nichtlinearen integralgleighungen. *Math. Annalen.*, 63:433–476, 1906.

[153] L. J. Schulman and U. Vazirani. Scalable NMR quantum computation. *arXive e-print quant-ph/9804060*, 1998.

[154] B. Schumacher. Quantum coding. *Phys. Rev. A*, 51:2738–2747, 1995.

[155] B. Schumacher and M. D. Westmoreland. Sending classical information via noisy quantum channels. *Physical Review A*, 56(1):131–138, 1997.

[156] B. Schumacher and M. D. Westmoreland. Quantum privacy and quantum coherence. *Phys. Rev. Lett.*, 80(25):5695–5697, 1998.

[157] B. W. Schumacher. Sending entanglement through noisy quantum channels. *Phys. Rev. A*, 54:2614, 1996.

[158] B. W. Schumacher and M. A. Nielsen. Quantum data processing and error correction. *Physical Review A*, 54(4):2629, 1996.

[159] B. W. Schumacher, M. Westmoreland, and W. K. Wootters. Limitation on the amount of acessible information in a quantum channel. *Phys. Rev. Lett.*, 76:3453, 1996.

[160] C. E. Shannon. A mathematical theory of communication. *Bell System Tech. J.*, 27:379–423, 623–656, 1948.

[161] C. E. Shannon. Communication theory of secrecy systems. *Bell Sys. Tech. J.*, 28:656–715, 1949.

[162] C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, Urbana, 1949.

[163] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings, 35th Annual Symposium on Fundamentals of Computer Science*, Los Alamitos, 1994. IEEE Press.

[164] P. W. Shor. Scheme for reducing decoherence in quantum memory. *Phys. Rev. A*, 52:2493, 1995.

[165] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comp.*, 26(5):1484–1509, 1997.

[166] P. W. Shor and J. A. Smolin. Quantum error-correcting codes need not completely reveal the error syndrome. *arXive e-print quant-ph/9604006*, 1996.

[167] G. F. Simmons. *Introduction to Topology and Modern Analysis*. McGraw-Hill, Auckland, 1963.

[168] B. Simon. *Trace ideals and their applications*. Cambridge University Press, Cambridge, 1979.

[169] D. T. Smithey, M. Beck, M. G. Raymer, and A. Faridani. Measurement of the Wigner distribution and the density-matrix of a light mode using optical homodyne tomography: application to squeezed states and the vacuum. *Phys. Re. Lett.*, 70(9):1244–1247, 1993.

[170] S. L. Sondhi, S. M. Girvin, J. P. Carini, and D. Shahar. Continuous quantum phase transitions. *Rev. Mod. Phys.*, 69(1):315–333, 1997.

[171] A. Steane. The ion-trap quantum information processor. *App. Phys. B – Lasers and Optics*, 64(6):623–642, 1997.

[172] A. M. Steane. Error correcting codes in quantum theory. *Physical Review Letters*, 77:793, 1996.

[173] B. M. Terhal and J. A. Smolin. Superfast quantum algorithms for coin weighing and binary search problems. *arXive e-print quant-ph/9705041*, 1997.

[174] The Quantum Computation Collective. What makes quantum computers powerful? 1997.

[175] K. S. Thorne. Gravitational waves: A new window onto the universe. In V. L. Fitch, D. R. Marlow, and M. A. E. Dementi, editors, *Critical problems in physics*, pages 167–197. Princeton University Press, Princeton, New Jersey, 1997.

[176] Q. A. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble. Measurement of conditional phase shifts for quantum logic. *Phys. Rev. Lett.*, 75:4710, 1995.

[177] A. M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proc. Lond. Math. Soc. 2 (reprinted in [56])*, 42:230, 1936.

[178] A. Uhlmann. The 'transition probability' in the state space of a *-algebra. *Reports on Mathematical Physics*, 9:273–279, 1976.

[179] T. van Gelder. The dynamical hypothesis in cognitive science. *To appear in Brain and Behavioral Sciences*, 1998.

[180] L. Viola and S. Lloyd. Dynamical suppression of decoherence in two-state quantum systems. *arXive e-print quant-ph/9803057*, 1998.

[181] K. Vogel and H. Risken. Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase. *Phys. Rev. A*, 40(12):7113–7120, 1989.

[182] A. Wehrl. General properties of entropy. *Rev. Mod. Phys.*, 50:221, 1978.

[183] S. Weinberg. *Dreams of a Final Theory.* Pantheon Books, 1993.

[184] S. Weinberg. *The Quantum Theory of Fields I.* Cambridge University Press, Cambridge, 1995.

[185] M. D. Westmoreland and B. Schumacher. Quantum entanglement and the non-existence of superluminal signals. *arXive e-print quant-ph/9801014*, 1998.

[186] J. A. Wheeler. Physics as meaning circuit. In G. E. Moore and M. O. Scully, editors, *Frontiers of nonequilibrium statistical physics*, volume 135 of *NATO ASI Series B*. Plenum Press, New York, 1986.

[187] J. A. Wheeler and W. H. Zurek. *Quantum theory and measurement.* Princeton University Press, Princeton NJ, 1983.

[188] S. Wiesner. Unpublished manuscript, circa 1969, appeared as [189].

[189] S. Wiesner. Conjugate coding. *SIGACT News*, 15:77, 1983.

[190] E. P. Wigner. Die messung quantenmechanischer operatoren. *Z. Phys.*, 133:101, 1952.

[191] F. Wilczek. The future of particle physics as a natural science. In V. L. Fitch, D. R. Marlow, and M. A. E. Dementi, editors, *Critical problems in physics*, pages 281–306. Princeton University Press, Princeton, New Jersey, 1997.

[192] C. P. Williams and S. H. Clearwater. *Explorations in quantum computing.* Springer-Verlag, New York, 1998.

[193] E. Witten. Vistas in theoretical physics. In V. L. Fitch, D. R. Marlow, and M. A. E. Dementi, editors, *Critical problems in physics*, pages 271–278. Princeton University Press, Princeton, New Jersey, 1997.

[194] S. Wolfram. *Cellular automata and complexity.* Addison-Wesley, Redwood City, CA, 1994.

[195] W. K. Wootters. Entanglement of formation of an arbitrary state of two qubits. *Phys. Rev. Lett.*, 80(10):2245–2248, 1998.

[196] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.

[197] A. C. Yao. Some complexity questions related to distributed computing. *Proc. of the 11th Ann. ACM Symp. on Theory of Computing*, pages 209–213, 1979.

[198] A. C. Yao. Quantum circuit complexity. *Proc. of the 34th Ann. IEEE Symp. on Foundations of Computer Science*, pages 352–361, 1993.

[199] A. Yariv. *Quantum electronics*. John Wiley and Sons, New York, 1989.

[200] C. Zalka. Simulating quantum systems on a quantum computer. *Proc. Roy. Soc.*, 454(1969):313–322, 1998.

[201] J. Ziv and A. Lempel. Compression of individual sequences by variable rate coding. *IEEE Trans. Inform. Theory*, IT-24:530, 1978.

[202] W. H. Zurek. Pointer basis of quantum apparatus: into what mixture does the wave packet collapse? *Phys. Rev. D*, 24(6):1516–1525, 1981.

[203] W. H. Zurek. Environment-induced super-selection rules. *Phys. Rev. D*, 26(8):1862–1880, 1982.

[204] W. H. Zurek. Algorithmic randomness and physical entropy. *Phys. Rev. A*, 40:4731, 1989.

[205] W. H. Zurek. Thermodynamic cost of computation, algorithmic complexity and the information metric. *Nature*, 341:119, 1989.

[206] W. H. Zurek. Decoherence and the transition from quantum to classical. *Phys. Tod.*, 44(10):36–44, 1991.

[207] W. H. Zurek. Preferred states, predictability, classicality and the environment-induced decoherence. *Prog. Theor. Phys.*, 89(2):281–312, 1993.

[208] W. H. Zurek. Decoherence, einselection, and the existential interpretation (the rough guide). *arXive e-print quant-ph/9805065*, 1998.

Index

L_1 distance
  classical, 90
χ, 112

absolute distance
  classical, 90
    operational meaning for, 91
  quantum, 92
algorithmic information, 182
amplitude damping, 61
angle, 99
Araki-Lieb inequality, 78

Bell basis, 22
binary entropy, 67
bit flip channel, 170
bit flip code, 171
Bloch sphere, 20
Bloch vector, 20

canonical representation for a quantum operation, 165
capacity
  theorem for noiseless qubit communication, 116
cavity QED, 31
chaining properties of distance measures, 104
channel capacity, 188
Church-Turing thesis, 6
  modern form, 27
classical physics, 224
clean protocols for quantum communication, 119
coherent information, 167, 192
communication complexity, 110
  coherent quantum, 110, 122
  of the inner product, 118
  quantum, 110
  unified model for, 126
complete positivity, 44
  example of a positive map not completely positive, 44
complete quantum operations, 43
computational basis states, 20
concavity of the entropy, 79
concavity of the quantum conditional entropy, 84
conditional entropy
  classical, 69
  quantum, 76
    concavity of, 84
continuity properties of distance measures, 104
controlled not gate, 24
convexity of quantum relative entropy, 84
correlation entropy, 180

data compression, 130
data pipelining inequality, 194
  classical, 73
data processing inequality
  classical, 72
  quantum, 167
decoherence, 224
discrete memoryless quantum channel, 199
distance measures, 89
dynamic distance, 101
dynamic error, 102
dynamic fidelity, 101

entanglement fidelity, 101
entropy
  classical, 66
  concavity of, 79
  of an ensemble, 77
  quantum, 73
entropy exchange, 164
entropy-fidelity lemma, 194
environmental models
  complete quantum operations, 48

incomplete quantum operations, 49
error, 98
error correction, 169

Fannes' inequality, 95
Fano inequality
classical, 115
quantum analogue to, 166
fidelity
classical, 90
quantum, 96
explicit formula, 97
fundamental physics, 13

generalized measurements, 62

Hadamard gate, 21
halting observable, 7
family of observables, 29
halting problem, 7
proof of unsolvability, 28
Hamming distance, 90
Holevo χ, 112
Holevo bound, 111

identity gate, 21
incomplete quantum operations, 43
information source
classical, 90
quantum, 136
information theory
motivation for definitions, 67
information-theoretic conditions for quantum error correction, 175
inner product, 118
ion trap, 31

joint entropy theorem, 75

Klein's inequality, 74

Landauer limit, 17, 182
law of large numbers, 135
Lieb's theorem, 80
local unitary operations, 205
logical zero and one, 171

magic basis, 153
Maxwell's demon, 181
meaning circuits, 15
measurement, 41
mixed state Schmidt decomposition, 229
Moore's law, 16
quantum corollary to, 17
mutual information
classical, 69
quantum, 76

NMR, 30, 31
noiseless coding theorem, 67
noisy channel coding theorem, 188
non-increasing property of χ, 112
not gate, 21
nuclear magnetic resonance, 30, 31

observed channel, 207
operator-sum representation, 45
freedom in, 51

partial trace property of χ, 112
Pauli sigma matrices, 172
Pauli sigmas matrices, 20
phase flip operator, 21, 172
phase shift operator, 21
phase transitions and entanglement, 161
physical quantum operations, 43
physics of information, 3
positive operator valued measures, 62
POVMs, 62
process tomography, 57
for incomplete quantum operations, 61
purification, 227

quantum channel capacity, 188
quantum circuit, 26
for teleportation, 24
requirements for, 31
quantum communication complexity, 110
of the inner product, 118
unified model for, 126
quantum computation, 26
requirements for, 31
quantum data compression, 130, 136
quantum data processing inequality, 167
quantum error correction, 169
quantum Fano inequality, 166
quantum gates, 21
quantum information, 8
quantum information processors, 30

quantum operations, 41, 43
  complete, 43
    environmental models for, 48
  definition, 44
  incomplete, 43
    environmental models for, 49
  limitations to the formalism, 63
  on a qubit, 50
  partial trace map, 47
  physical, 43
  trace map, 47
quantum process tomography, 57
  for incomplete quantum operations, 61
quantum source, 198
quantum state tomography, 57
quantum teleportation
  connection to error correction, 56
  experimental implementation, 36
qubit, 19
  quantum operations on, 50

relative entropy
  classical, 68
  quantum, 74
    convexity of, 84
    strengthened convexity result, 88
reversible computation, 17

Schmidt bases, 131, 228
Schmidt decomposition, 131, 227
  for mixed states, 229
Schmidt number, 131, 228
Schumacher's theorem, 130
Shannon entropy, 66
Shannon's noiseless coding theorem, 67
Shor's quantum error correcting code, 173
sigma matrices, 20
state tomography, 57
strengthened convexity of quantum relative entropy, 88
strong subadditivity
  classical, 70
  quantum, 80
    proof of, 85
subadditivity, 196
  classical, 70
  of mutual information, 72
  of the quantum conditional entropy, 86
  quantum, 78
superdense coding, 21

TCE, 32
teleportation, 23
  as a quantum operation, 54
tomography, 57
transpose operation, 44
triangle inequality, 78
  for communication complexity, 127
trichloroethylene, 32
typical sequence, 133
typical subspace, 133

Uhlmann's formula for fidelity, 96
uniformity requirement for quantum computation, 27
unifying picture for quantum information, 222
universal set of gates, 21

von Neumann entropy, 73

weak law of large numbers, 135